HP StorageWorks

P4000 G2 Unified NAS Gateway User Guide

Part Number: 5697-0624
First edition: November 2010

Legal and notice information

Acknowledgements

Microsoft, Windows, Windows XP, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

Warranty

WARRANTY STATEMENT: To obtain a copy of the warranty for this product, see the warranty information website:

http://www.hp.com/go/storagewarranty

Contents

Figures

Tables

1 Component identification

This chapter provides illustrations of the storage system hardware components.

NOTE:

The keyboard, mouse, and monitor are used only for the direct attached method of accessing the server. They are not provided with your storage system.

P4000 G2 Unified NAS Gateway hardware components

The following figures show components and LEDs located on the front and rear panels of the P4000 G2 Unified NAS Gateway.

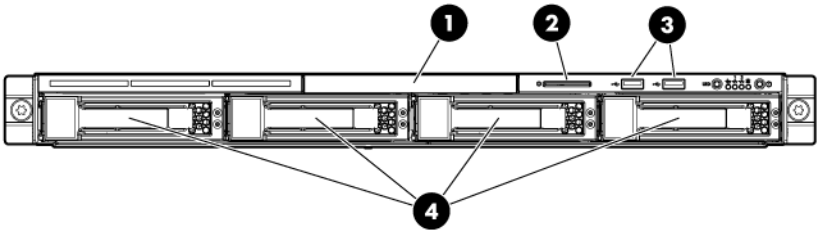

Figure 1 P4000 G2 Unified NAS Gateway front panel components

1. DVD-RW drive
2. Serial label pull tab
3. Two (2) USB ports
4. Four (4) 3.5″ hot-plug SAS/SATA hard drive bays

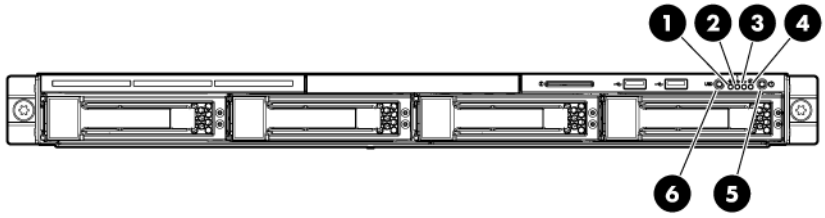

Figure 2 P4000 G2 Unified NAS Gateway front panel LEDs

Table 1 P4000 G2 Unified NAS Gateway front panel LED descriptions

Item / Description	Status
1. Internal health LED	Green = System health is normal. Amber = System health is degraded. Red = System health is critical. Off = System health is normal (when in standby mode).
2. NIC 1 link/activity LED 3. NIC 2 link/activity LED	Green = Network link exists. Flashing green = Network link and activity exist. Off = No network link exists.
4. Drive activity LED	Green = Drive activity is normal. Off = No drive activity exists.
5. Power On/Standby button and system power LED	Green = Normal (system on) Amber = System is in standby, but power is still applied. Off = Power cord is not attached or the power supply has failed.
6. UID button/LED	Blue = Identification is activated. Flashing blue = System is being managed remotely. Off = Identification is deactivated.

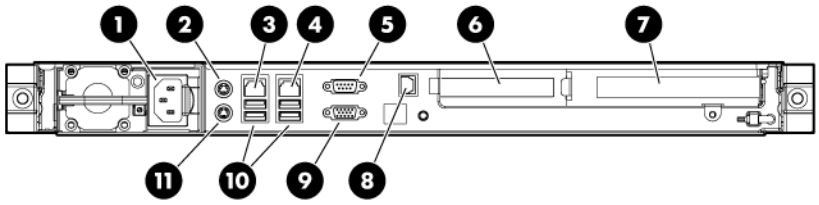

Figure 3 P4000 G2 Unified NAS Gateway rear panel components

1. Power cord connector
2. Mouse connector
3. 10/100/1000 NIC 1 connector/shared iLO 2 management port
4. 10/100/1000 NIC 2 connector
5. Serial connector
6. Low profile PCIe slot (occupied by Smart Array P212 controller)
7. Full-sized PCIe slot (occupied by NC364T 4-port NIC)
8. Dedicated iLO 2 management port (this port is optional and must be purchased separately)
9. Video connector
10. USB connectors (2)
11. Keyboard connector

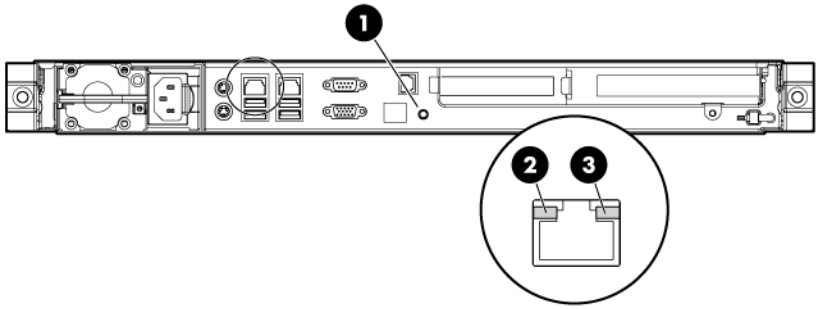

Figure 4 P4000 G2 Unified NAS Gateway rear panel LEDs

Table 2 P4000 G2 Unified NAS Gateway rear panel LED descriptions

Item / Description	Status
1. UID button/LED	Blue = Activated Flashing = System is being managed remotely. Off = Deactivated
2. NIC/iLO 2 link	Green or flashing green = Activity exists. Off = No activity exists.
3. NIC/iLO 2 activity	Green = Link exists. Off = No link exists.

2 Installing and configuring the server

Setup overview

The HP StorageWorks P4000 G2 Unified NAS Gateway comes preinstalled with the Microsoft Windows® Storage Server™ 2008 R2 Enterprise x64 Edition operating system with Microsoft iSCSI Software Target and a Microsoft Cluster Service (MSCS) license included.

IMPORTANT:

- Windows Storage Server 2008 R2 x64 operating systems are designed to support 32–bit applications without modification; however, any 32–bit applications that are run on these operating systems should be thoroughly tested before releasing the storage system to a production environment.
- Windows Storage Server x64 editions support only x64-based versions of Microsoft Management Console (MMC) snap-ins, not 32-bit versions.

Determine an access method

Before you install the storage system, you need to decide on an access method.

The type of access you select is determined by whether or not the network has a Dynamic Host Configuration Protocol (DHCP) server. If the network has a DHCP server, you can install the storage system through the direct attachment or remote management methods. If your network does not have a DHCP server, you must access the storage system through the direct attachment method.

The direct attachment method requires a display, keyboard, and mouse. These components are not provided with the storage system.

IMPORTANT:

Only the direct attach and remote management access methods can be used to install the storage system. After the storage system installation process is complete and the system's IP address has been assigned, you can then additionally use the remote browser and remote desktop methods to access the storage system.

Check kit contents

Remove the contents, making sure you have all the components listed below. If components are missing, contact HP technical support.

- HP StorageWorks P4000 G2 Unified NAS Gateway (with operating system preloaded)
- Power cord(s)
- Product Documentation and Safety and Disposal Documentation CD

- *HP StorageWorks Storage System Recovery DVD*
- *End User License Agreement*
- Certificate of Authenticity Card
- Slide rail assembly
- *HP ProLiant Essentials Integrated Lights-Out 2 Advanced Pack*

Locate the serial number, Certificate of Authenticity, and End User License Agreement

For technical support purposes, locate the storage system's serial number, Certificate of Authenticity (COA), and End User License Agreement (EULA). Record the serial number and COA product key and make a print copy of the EULA as needed.

The storage system's serial number is located in several places:

- Top of the storage system
- Back of the storage system
- Inside the storage system shipping box
- Outside of the storage system shipping box

The storage system's Certificate of Authenticity (COA) card is located inside the storage system shipping box. There is also a COA sticker with product key affixed to the top of the storage system.

The storage system's printed End User License Agreement (EULA) is located in the media kit that is shipped with the storage system. There is also an electronic copy of the EULA installed with the storage system at **%SystemDrive%\Windows\System32\license.rtf**.

Install the storage system hardware

1. Install the rail kit by following the *HP Rack Rail Kit installation instructions*.
2. If connecting to the storage system using the direct attach method, connect the following cables to the back panel of the storage system in the following sequence: keyboard, mouse, network cable, monitor cable, and power cable.

NOTE:

- The keyboard, mouse, and monitor are not provided with the storage system.

3. If connecting to the storage system using the remote management method, connect a network cable to a data port, a network cable to the iLO 2 port, and power cable.

Connect to the storage system

Use either the direct connect or remote management method to connect to the storage system.

IMPORTANT:

Only the direct attach and remote management access methods can be used to install the storage system. After the storage system installation process is complete and the system's IP address has been assigned, you can then additionally use the remote browser and remote desktop methods to access the storage system.

- Direct attach — Connect the following cables to the back panel of the storage system in the following sequence: keyboard, mouse, network cable, monitor cable, and power cable. This access method is mandatory if your network does not have a Dynamic Host Configuration Protocol (DHCP) server.

NOTE:

The keyboard, mouse, and monitor are not provided with the storage system.

- Remote management — Access the storage system using the Integrated Lights-Out remote management method:
  1. Ensure that a network cable is connected to the iLO port located on the back of the storage system.
  2. Locate the iLO Network Settings tag attached to the storage system and record the default user name, password, and DNS name.
  3. From a remote computer, open a standard Web browser and enter the iLO management hostname of the storage system.

     NOTE:

     By default, iLO obtains the management IP address and subnet mask from your network's DHCP server. The hostname found on the iLO tag is automatically registered with your network's DNS server.

  4. Using the default user information provided on the iLO Network Settings tag, log on to the storage system.

For detailed instructions on using iLO remote management software, see the *HP Integrated Lights–Out 2 User Guide* or *HP ProLiant Integrated Lights-Out 3 User Guide*.

Power on the server and log on

Power on the server after installing the hardware and connecting the cables. Powering on the server for the first time initiates the storage system installation process.

1. Power on the system by pushing the power button on the front panel. If using iLO 2, click **Momentary Press** on the **Power Management** page to power on the server, then click **Launch** on the **Status Summary** page to open the iLO 2 Integrated Remote Console and complete the installation process.

   The storage system starts and displays an HP Network Storage System installation screen. The storage system installation process takes approximately 10–15 minutes.

NOTE:

Your storage system comes pre-installed with the Microsoft Windows Storage Server 2008 R2 operating system. There is no operating system installation required.

When the storage system installation process nears completion, the Windows Storage Server 2008 R2 desktop displays the following message: **The user's password must be changed before logging on the first time**. Log on to the storage system by establishing an Administrator password:

2. Click **OK**.
3. Type an Administrator password in the **New password** box.
4. Re-type the Administrator password in the **Confirm password** box.
5. Click the blue arrow next to the **Confirm password** box.
6. Click **OK**.

   After the Administrator password has been set, the storage system completes the installation process and restarts.
7. When prompted, press `CTRL+ALT+DELETE` to log on to the system. If using iLO 2, on the iLO 2 Integrated Remote Console tab, click the button labeled **CAD** and then click the **Ctrl-Alt-Del** menu item.

IMPORTANT:

After establishing the new Administrator password, be sure to remember it and record it in a safe place if needed. HP has no way of accessing the system if the new password is lost.

After logging in for the first time, the Welcome screen of the HP StorageWorks Rapid Startup Wizard opens. Use the HP StorageWorks Rapid Startup Wizard to set up your system with basic configuration information.

Confirm Windows activation

Immediately after installing the storage system, confirm that your copy of the Microsoft® Windows® Storage Server™ 2008 R2 operating system has been activated.

IMPORTANT:

Some storage systems may not automatically activate Windows during the system installation process.

To check the activation status of Windows, open Control Panel (**Start > Control Panel**) and then double-click **System**. If your copy of Windows has been activated, the status under **Windows activation**

will read **Windows is activated**. If your copy of Windows has not been activated, you must manually activate it within three days of the initial storage system installation.

To manually activate your copy of Windows:

1. If needed, change your locale settings (**Control Panel > Regional and Language Options**).
2. Record the product key located on the Certificate of Authenticity.
3. In Control Panel, double-click **System**.
4. Under **Windows Activation**, click **Activate Windows Now**.
5. Follow the on-screen instructions. Be sure to choose the **Type a different product key** option when prompted for a product key.

Microsoft Windows Storage Server 2008 R2 cannot be fully activated online. During the activation process, you will be provided with a telephone number to call to complete the process. Be sure to have the product key available when calling this number.

Configure the storage system

The **HP StorageWorks Initial Configuration Tasks** window launches automatically on logon. Use this tool to perform setup tasks such as setting the time zone, network configuration, changing the computer name, and joining a domain. When the HP StorageWorks Initial Configuration Tasks window is closed, **HP Server Manager** is launched automatically. Use HP Server Manager for further customizing of the storage system, such as adding roles and features, and share and storage management.

For detailed information about configuration options in HP StorageWorks Initial Configuration Tasks and Server Manager, see the online help.

Complete system configuration

After the storage system is physically set up and the basic configuration is established, you must complete additional setup tasks. Depending on the deployment scenario, these steps can vary. Additional steps may include:

- Running Microsoft Windows Update — HP highly recommends running Microsoft Windows updates to identify, review, and install the latest, applicable, critical security updates.
- Creating and managing users and groups—User and group information and permissions determine whether a user can access files. If the storage system is deployed into a workgroup environment, this user and group information is stored locally on the device. By contrast, if the storage system is deployed into a domain environment, user and group information is stored on the domain.
- Joining workgroups and domains—These are the two system environments for users and groups. Because users and groups in a domain environment are managed through standard Windows or Active Directory domain administration methods, this document discusses only local users and groups, which are stored and managed on the storage system. For information on managing users and groups on a domain, see the domain documentation available on the Microsoft web site.

  If the storage system is deployed in a domain environment, the domain controller will store new accounts on the domain; however, remote systems will store new accounts locally unless they are granted permissions to create accounts on the domain.
- Using Ethernet NIC teaming (optional)—Select models are equipped with an HP or Broadcom NIC Teaming utility. The utility allows administrators to configure and monitor Ethernet network interface controller (NIC) teams in a Windows-based operating system. These teams provide options for increasing fault tolerance and throughput.

- Activating iLO 2 Advanced features using a license key—The Remote Console feature of iLO 2 requires a license key. The key is included with the storage system inside the Country Kit. See the iLO 2 Advanced License Pack for activation instructions.
- Adjusting logging settings for system, application, and security events.
- Installing third-party software applications such as an antivirus application.
- Registering the server — To register the server, refer to the HP Registration website (http://register.hp.com).

Additional access methods

After the storage system installation process is complete and the system's IP address has been assigned, you can then additionally use the remote browser, Remote Desktop, and Telnet Server methods to access the storage system.

Using the remote browser method

The storage system ships with DHCP enabled on the network port. If the server is placed on a DHCP-enabled network and the IP address or server name is known, the server can be accessed through a client running Internet Explorer 5.5 (or later) on that network, using the TCP/IP 3202 port.

IMPORTANT:

Ensure that you have the following:

- Windows-based PC loaded with Internet Explorer 5.5 (or later) on the same local network as the storage system
- DHCP-enabled network
- Server name or IP address of the storage system

To connect the server to a network using the remote browser method, ensure that the client is configured to download signed ActiveX controls.

To connect the storage system to a network using the remote browser method

1. On the remote client machine open Internet Explorer and enter `https://` and the server name of the storage system followed by a hyphen (-), and then `:3202`. For example, `https://labserver-:3202`.

NOTE:

If known, you can substitute the IP address for the server name. For example: `192.100.0.1:3202`.

2. Click **OK** on the **Security Alert** prompt.
3. When prompted, log on to the storage system with the administrator user name and password.

IMPORTANT:

When using the remote browser method to access the storage system, always close the remote session before closing your Internet browser. Closing the Internet browser does not close the remote session. Failure to close your remote session impacts the limited number of remote sessions allowed on the storage system at any given time.

Using the Remote Desktop method

Remote Desktop provides the ability for you to log onto and remotely administer your server, giving you a method of managing it from any client. Installed for remote administration, Remote Desktop allows only two concurrent sessions.

To connect the storage system to a network using the Remote Desktop method

1. On the PC client, select **Start > Run**. At **Open**, type `mstsc`, then click **OK**.
2. Enter the IP address of the storage system in the **Computer** box and click **Connect**.
3. When prompted, log on to the storage system with the administrator user name and password.

Using the Telnet method

Telnet is a utility that lets you connect to servers, log on, and obtain a command prompt remotely. Telnet is included with the OS but must be activated before use.

CAUTION:

For security reasons, Telnet is disabled by default. The service needs to be modified to enable access to the storage system with Telnet.

Enabling Telnet

1. In Server Manager, expand the **Configuration** node in the left panel.
2. Click **System and Network Settings**.
3. Under **System Settings Configuration**, click **Telnet**.
4. Select **Enable Telnet access to this server** and then click **OK**.

Default storage settings

HP StorageWorks P4000 G2 Unified NAS Gateway is preconfigured with default storage settings. This section provides additional details about the preconfigured storage.

Physical configuration

The logical disks reside on physical drives as shown in the table below.

IMPORTANT:

The first two logical drives are configured for the storage system operating system.

The Operating System volume default factory settings can be customized after the operating system is up and running. The following settings can be changed:

- RAID level can be changed to any RAID level except RAID 0
- OS logical drive size can be changed to 40 GB or higher

If the Operating System volume is customized and the System Recovery DVD is run at a later time, the System Recovery process will maintain the custom settings as long as the above criteria are met (RAID level other than RAID 0 and OS logical drive size of 40 GB or higher) and the OS volume is labeled **System**. If the storage system arrays are deleted and the System Recovery DVD is run, the System Recovery process will configure the storage system using the factory default settings listed in the table below.

HP StorageWorks P4000 G2 Unified NAS Gateways do not include preconfigured data volumes. The administrator must configure data storage for the storage system. See "Configuring data storage" on page 68 for more information.

The system reserved partition contains the operating system boot loader and allows you to enable BitLocker Drive Encryption for the Operating System volume.

Table 3 P4000 G2 Unified NAS Gateway RAID configuration

Model	Logical Disk 1
HP StorageWorks P4000 G2 Unified NAS Gateway	• Operating System Volume • RAID 1+0 • Physical Drives 0–1

NOTE:

In the HP Array Configuration Utility (ACU), logical disks are labeled 1 and 2. In Microsoft Disk Manager, logical disks are displayed as 0 and 1. For HP Smart Array configuration information, see http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/.

If the operating system has a failure that might result from corrupt system files, a corrupt registry, or the system hangs during boot, see "System recovery" on page 105.

Default boot sequence

The BIOS supports the following default boot sequence:

1. DVD-ROM
2. HDD
3. Bootable USB flash drive
4. PXE (network boot)

Under normal circumstances, the storage systems boot up from the OS logical drive.

- If the system experiences a drive failure, the drive displays an amber disk failure LED.

- If a single drive failure occurs, it is transparent to the OS.

3 Cluster configuration

HP StorageWorks P4000 G2 Unified NAS Gateway hardware components are configured in a clustered environment to a P4000 SAN Solution. The P4000 SAN Solution provides iSCSI block services while the P4000 G2 Unified NAS Gateway provides data services using CIFS/NFS protocols. The P4000 G2 Unified NAS Gateway also supports optional data protection services with tools such as HP Data Protector, Data Protection Manager (DPM), VMWare, and VMware Consolidated Backup (VCB).

The P4000 G2 Unified NAS Gateways are clustered and connected to the HP P4000 SAN network segment using standard Ethernet IP switches. Up to eight (8) P4000 G2 Unified NAS Gateway nodes can be clustered together.

The following figure shows a complete network infrastructure comprised of two P4000 G2 Unified NAS Gateways connected to a P4000 SAN segment using standard IP switches while also connected to the client network using a standard IP switch. Application servers are also connected the P4000 SAN segment and client network using standard IP switches.

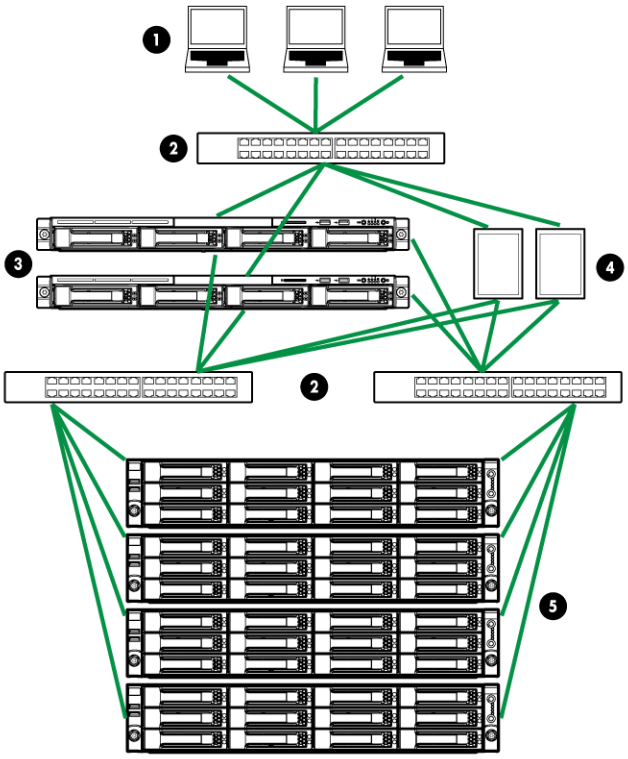

Figure 5 P4000 G2 Unified NAS Gateway network infrastructure

1. Client network
2. Standard IP switch
3. P4000 G2 Unified NAS Gateway nodes
4. Application servers

5. P4000 SAN segment

IMPORTANT:

Instructions and illustrations in this document describe the installation and configuration of a 2–node P4000 G2 Unified NAS Gateway. If you purchased the 1–node P4000 G2 Unified NAS Gateway, all instructions related to installing and configuring the second node of the solution do not apply. The 1–node solution does not support full High Availability (HA) capability, but is HA ready and can be upgraded to a full HA solution by purchasing and installing an additional 1–node P4000 G2 Unified NAS Gateway.

After installing the P4000 G2 Unified NAS Gateway nodes as detailed in the *HP StorageWorks P4000 G2 Unified NAS GatewayQuick Start Guide*, the system components should be racked, cabled, powered on, and you should be logged in to the systems with Administrative privileges. In addition, P4000 SAN storage should be created and mapped to the P4000 G2 Unified NAS Gateway nodes using P4000 Centralized Management Console (CMC) software, Microsoft iSCSI Initiator, and Windows Disk Manager.

For complete information about creating and managing HP StorageWorks P4000 SAN Solutions, see the user documentation at http://www.hp.com/go/p4000. Click **HP Support and Drivers**, select your HP StorageWorks P4000 SAN Solution model, and then click **Manuals**.

Creating and configuring the cluster

The following section describes the process of creating the P4000 G2 Unified NAS Gateway cluster.

Set IP addresses for the network connections

In order to accurately describe the physical connections between the clustered components, the P4000 G2 Unified NAS Gateways are designated as **Server 1** and **Server 2** in this section.

Each P4000 G2 Unified NAS Gateway includes six NIC connectors, each reserved for a specific network connection purpose as shown in the figure below:

- One NIC connector is reserved for heartbeat connectivity between the two P4000 G2 Unified NAS Gateway storage systems (shown in red).
- Two NIC connectors are reserved for multi-path connectivity to P4000 SAN storage nodes (shown in green and blue).
- Three NIC connectors are reserved for front file serving into a client network (shown in black).

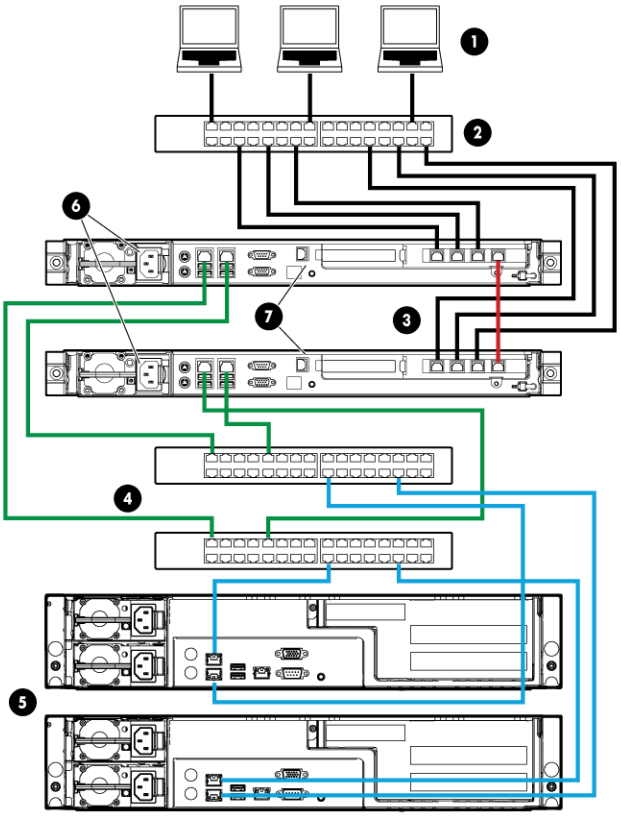

Figure 6 P4000 G2 Unified NAS Gateway connections

1. Client network
2. Standard IP switch
3. P4000 G2 Unified NAS Gateway nodes
4. Standard IP switches
5. P4000 SAN segment

For proper operation of the cluster, each storage server requires the following: a private heartbeat network connection between the two servers, at least one private network connection to a P4000 SAN, and at least one connection for file serving purposes, which can be configured as a public or private network connection based on your network infrastructure needs. The private connection NIC adapters must be set with static IP addresses; the public NIC adapters can be set with a static IP address, or may be automatically configured using DHCP. If a DHCP server is available on your network, HP recommends allowing DHCP to automatically configure the public-facing network connections; this is the default setting.

1. Log in to the Server 1 desktop as a user with Administrative privileges.
2. Click **Close** to dismiss the HP StorageWorks Rapid Startup Wizard.
3. In Server Manager, click **View Network Connections**.

   If Server Manager is not already open, click **Start > Administrative Tools > Server Manager**.

4. Identify the public and private connections:

   a. Right-click one of the connections and select **Status**.

      The connection status of the private connections will indicate **Local** in the **IPv4 Connectivity** field; the connection status of the public-facing connections will indicate **Internet** in this field.

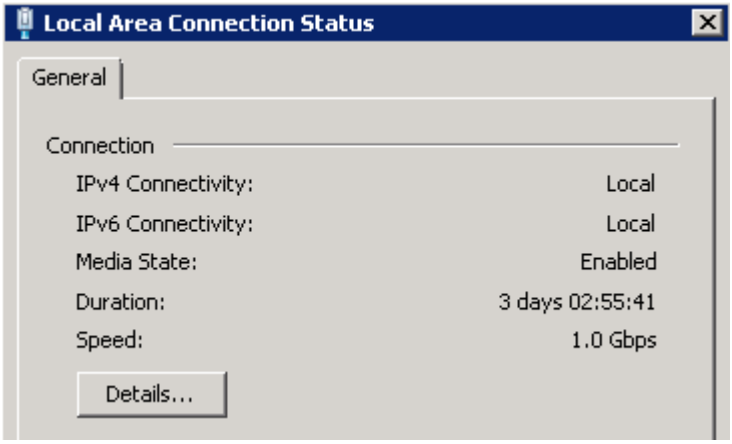

Figure 7 Private connection status

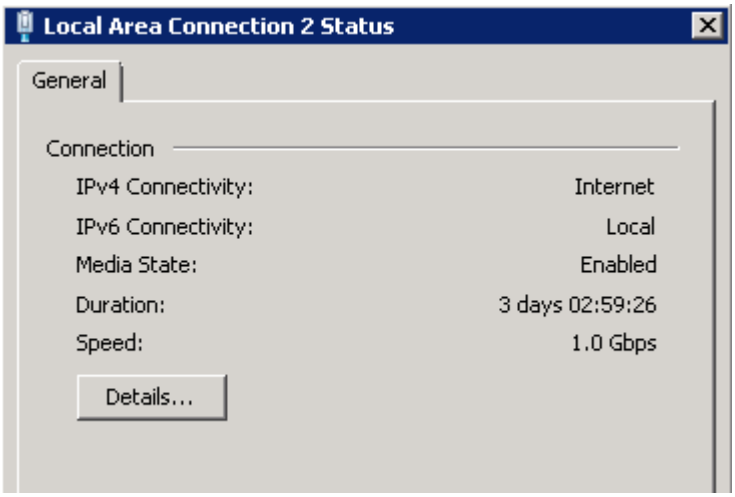

Figure 8 Public connection status

   b. After identifying the private and public connections, click **Close**.

TIP:

To more easily identify public and private network connections, rename them (for example, **Cluster Heartbeat**, **Public File Serve 1**, and **P4000 SAN 1**).

5. To assign static IP addresses to a private connection:
   a. Right-click the private connection and select **Properties**.
   b. Clear all items on the **General** tab except for **HP Network Configuration Utility**, **Internet Protocol Version 4 (TCP/IPv4)**, and **Internet Protocol Version 6 (TCP/IPv)**.
   c. Select **Internet Protocol Version 4 (TCP/IPv4)** and then click **Properties**.

   NOTE:

   For the purposes of this document, the IPv4 Internet Protocol is the documented IP version. If you are familiar with IPv6 and prefer to use it in your network environment, it is also supported.

   d. Select **Use the following IP address** and enter a static IP address and subnet mask using configuration information assigned by your network administrator.
   e. Click **Advanced**, select the **DNS** tab, and clear the **Register this connection's addresses in DNS** box.
   f. Click **OK** twice and then click **OK** to dismiss the **Local Area Connection Properties** dialog box.

   The Server 1 private static IP addresses are now set.
6. To set the Server 1 public IP addresses, do one of the following:
   - If a DHCP server is available on your network, allow DHCP to automatically configure your public-facing network connections.
   - If a DHCP server is not available on your network, configure a static IP address for the public-facing network connections as documented above using configuration information assigned by your network administrator. For the public-facing static IP address, do not clear any items on the **General** tab of the connection's **Properties** page.
7. Repeat the preceding steps on Server 2, setting the private and public IP addresses as needed.

Join both storage servers to the domain

1. From Server 1, open Server Manager, and click **Change System Properties**.

2. On the **Computer Name** tab, click **Change**.

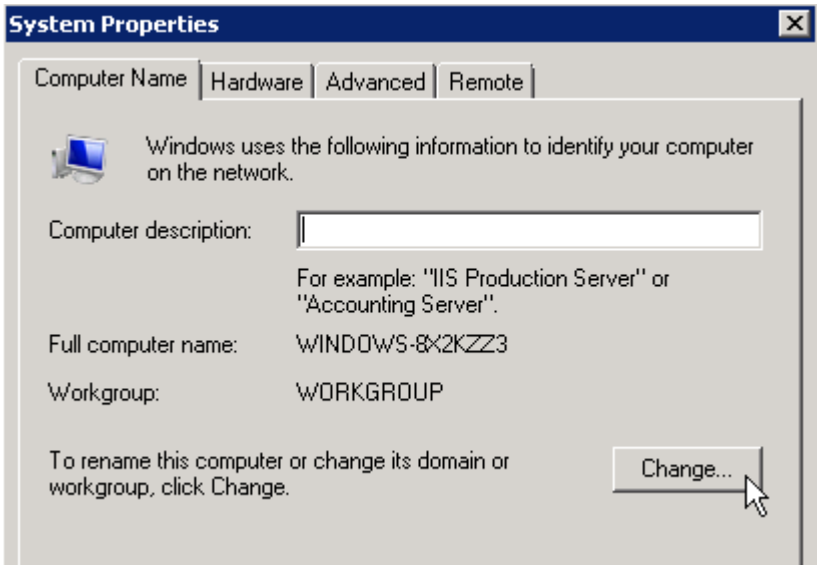

Figure 9 Computer Name tab of System Properties

3. On the **Computer Name/Domain Changes** dialog box, in the **Computer name** field, enter a unique name for the server.
4. Select the **Domain** radio button and type the name of the domain on which the cluster will reside and then click **OK**.

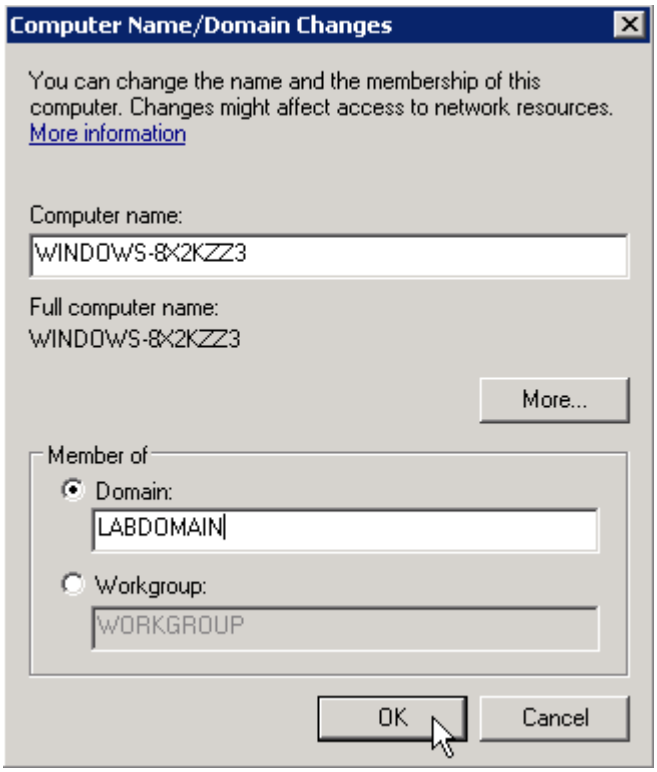

Figure 10 Computer Name Changes dialog box

5. When prompted for credentials, enter valid domain account credentials and then click **OK**.
6. Click **OK** to accept the domain changes.

7. When prompted, click **Yes** to restart the server.
8. Repeat these steps for Server 2.

Initialize and format the storage disks

The storage referenced in this section must be created and configured on the P4000 SAN. See the *HP StorageWorks P4000 Configuration Guide* for more information about connecting the SAN volumes to the Windows Storage Server instance.

1. From Server 1, open Server Manager, and under **Storage**, click **Disk Management**.
2. Right-click the **Disk 1** label and select **Online** to bring the disk online.
3. Right-click the disk, and then click **Initialize Disk**.

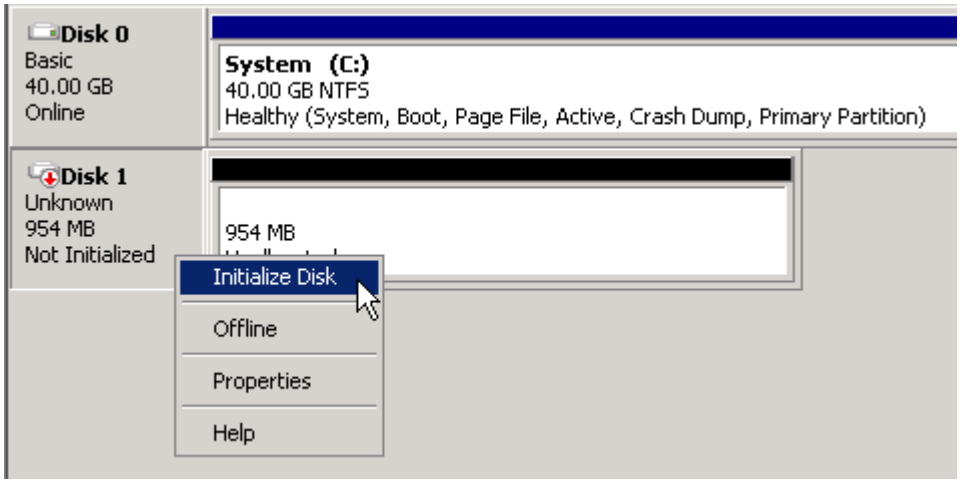

Figure 11 Initialize Disk 1 (the witness disk)

4. In the **Initialize Disk** dialog box, select the disk to initialize, select a partition style, and then click **OK**.

   The disk is initialized as a basic disk.
5. In the storage allocation area, right-click and select **New Simple Volume**.

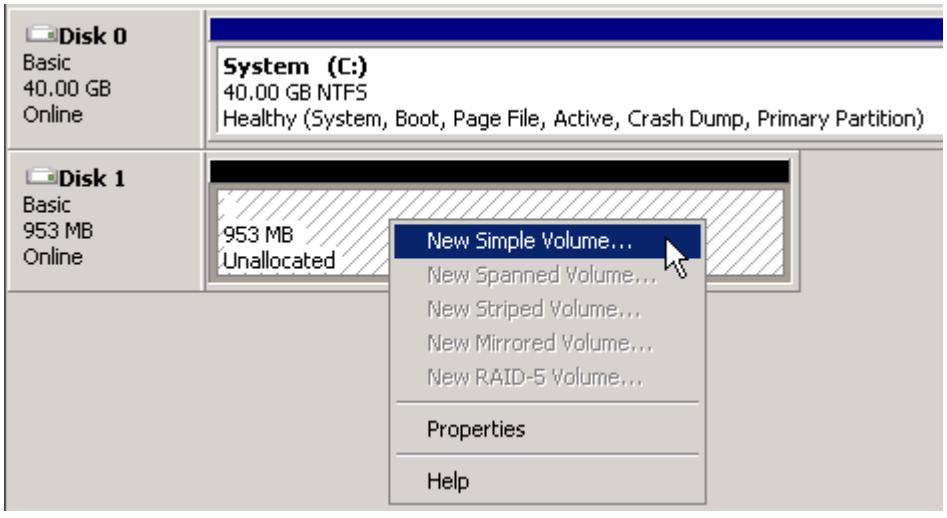

Figure 12 Create new simple volume

6. Complete the **New Simple Volume Wizard** with the following settings:
   - Accept the default assigned partition size
   - Assign drive letter **Q**
   - Formatted as **NTFS**
   - Label the volume **Witness**
   - Check **Perform a quick format**
7. Repeat the preceding steps to initialize and format the remaining data disks, assigning properties such as volume size and labels as appropriate for the intended use of the storage.

NOTE:

Before proceeding to additional configuration tasks, ensure that all disks have been completely initialized and formatted. After the disks have been completely initialized and formatted, the storage allocation area will indicate the volume name, size, and state (**Healthy**, for example).

Validate the configuration

The process of validating your configuration may take a few minutes. If additional storage is configured to be used by the cluster, the validation process takes additional time to complete.

1. From Server 1, click **Start > Administrative Tools > Failover Cluster Management**.

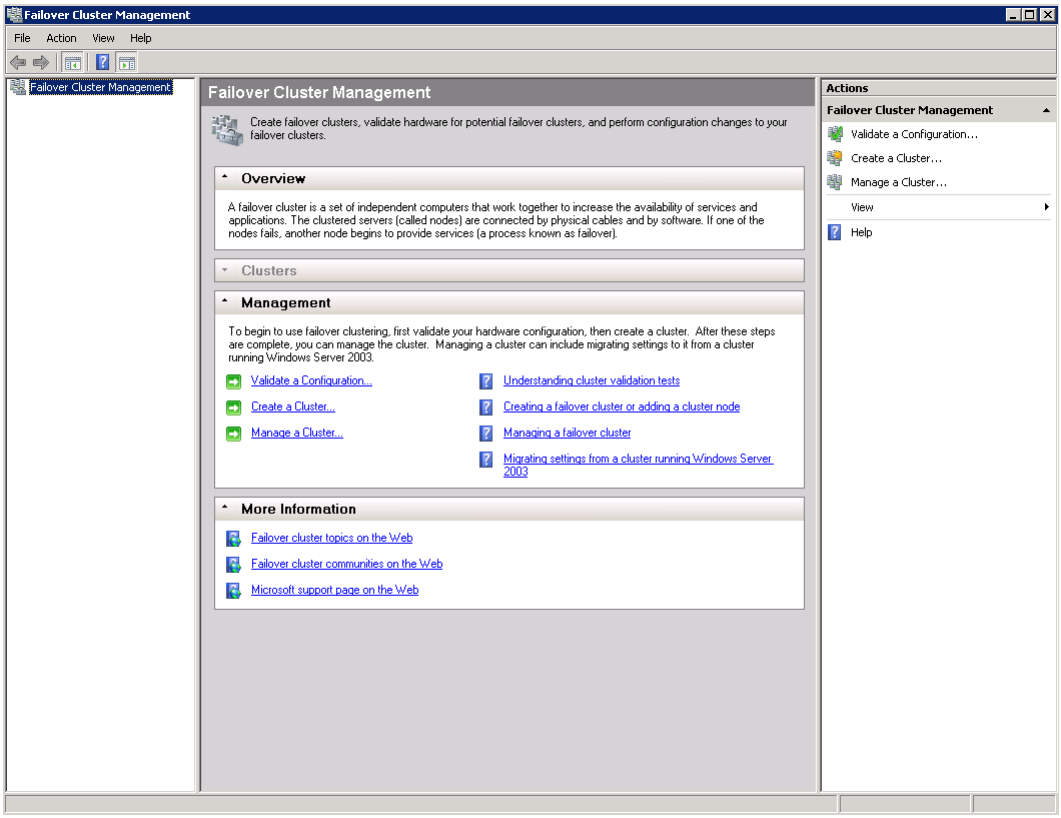

Figure 13 Failover Cluster Management user interface

2. In the **Management** section, click **Validate a Configuration**.

3. Click **Next**.
4. On the **Select Servers or a Cluster** page, enter and add the names of Server 1 and Server 2 and then click **Next**.

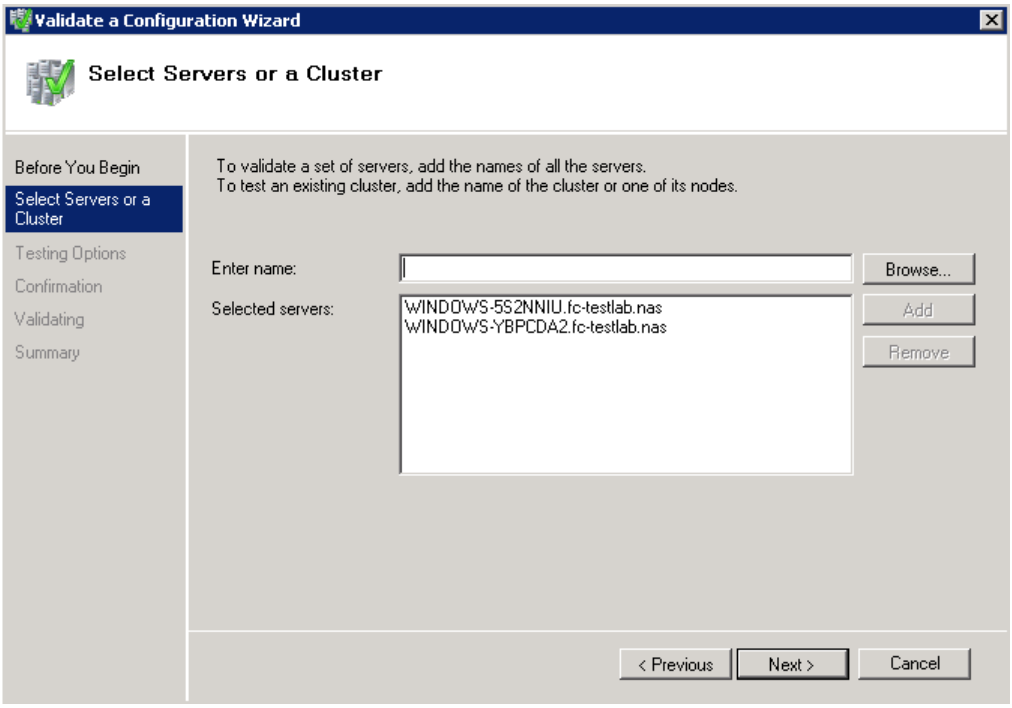

Figure 14 Select servers to be validated for the cluster

5. Select **Run all tests** and then click **Next**.

6. Review the details of the **Confirmation** page and then click **Next**.

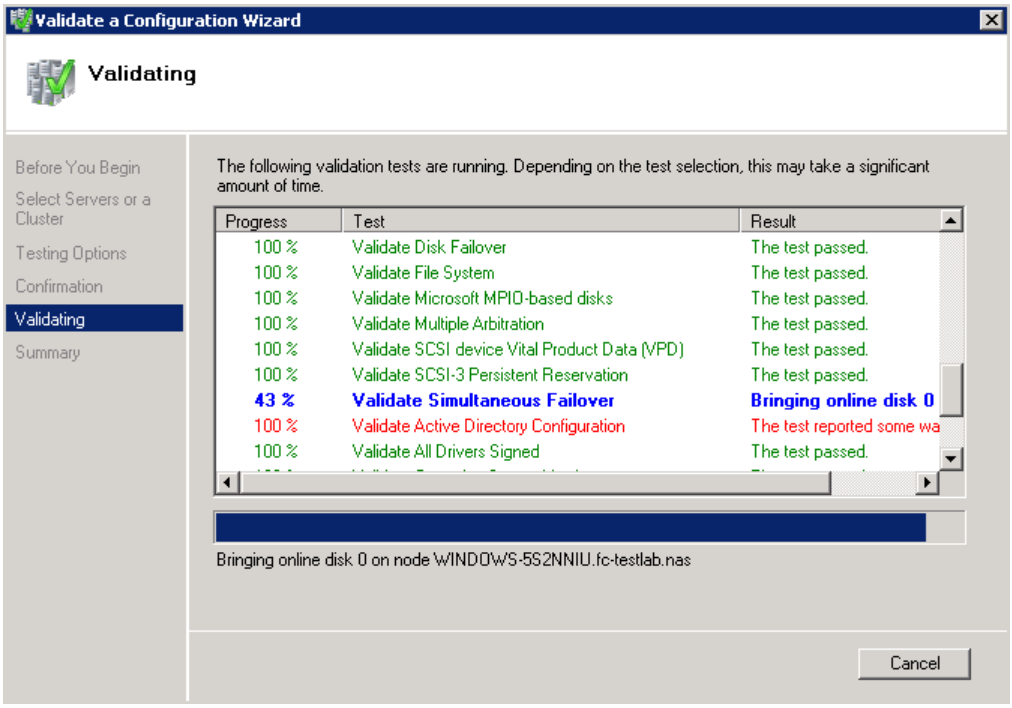

Figure 15 Validating the cluster configuration

7. After the validation tests have run, click **View Report** to review the validation test results.

   Use the information provided in the Failover Cluster Validation Report to troubleshoot issues that would prevent the successful creation of the cluster. After addressing the issues, re-run the **Validate a Configuration** wizard.

8. Click **Finish** to exit the **Validate a Configuration** wizard.

Create the cluster

1. In the **Failover Cluster Management** user interface, under **Management**, click **Create a Cluster**.
2. Click **Next**.
3. On the **Select Servers** page, enter and add the names of Server 1 and Server 2 and then click **Next**.

4. On the **Access Point for Administering the Cluster** page, type a unique name for the cluster and then click **Next**.

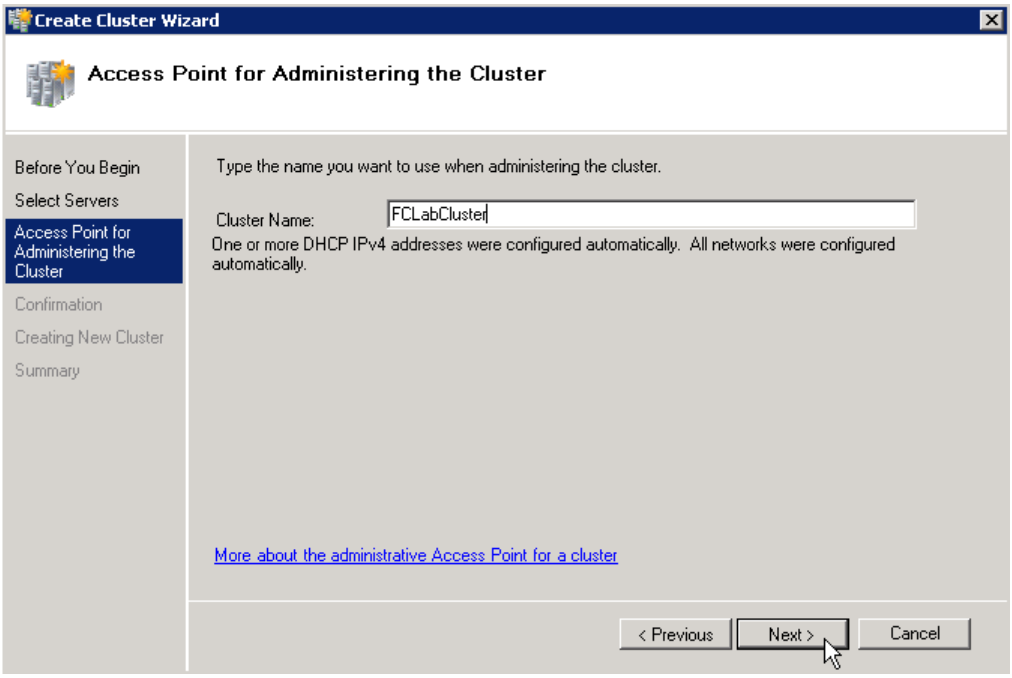

Figure 16 Entering cluster name

5. Review the information on the **Confirmation** page and then click **Next**.

   After the cluster is successfully created, the **Summary** page lists basic cluster information. Click **View Report** to view a detailed report. Click **Finish** to exit the **Create Cluster Wizard**.

Add services or applications to the cluster

For the purposes of this document, the following procedure illustrates adding a File Server; however, any supported service or application can be added for high availability. For example, you may choose to add NFS/CIFS or iSCSI LUNs to the cluster.

1. In the **Failover Cluster Management** user interface, locate the newly-created cluster in the left pane and then click **Services and Applications**.
2. In the **Actions** pane, click **Configure a Service or Application**.

   The **High Availability Wizard** appears.
3. On the **Before You Begin** page, click **Next**.

4. Select **File Server** from the list and then click **Next**.

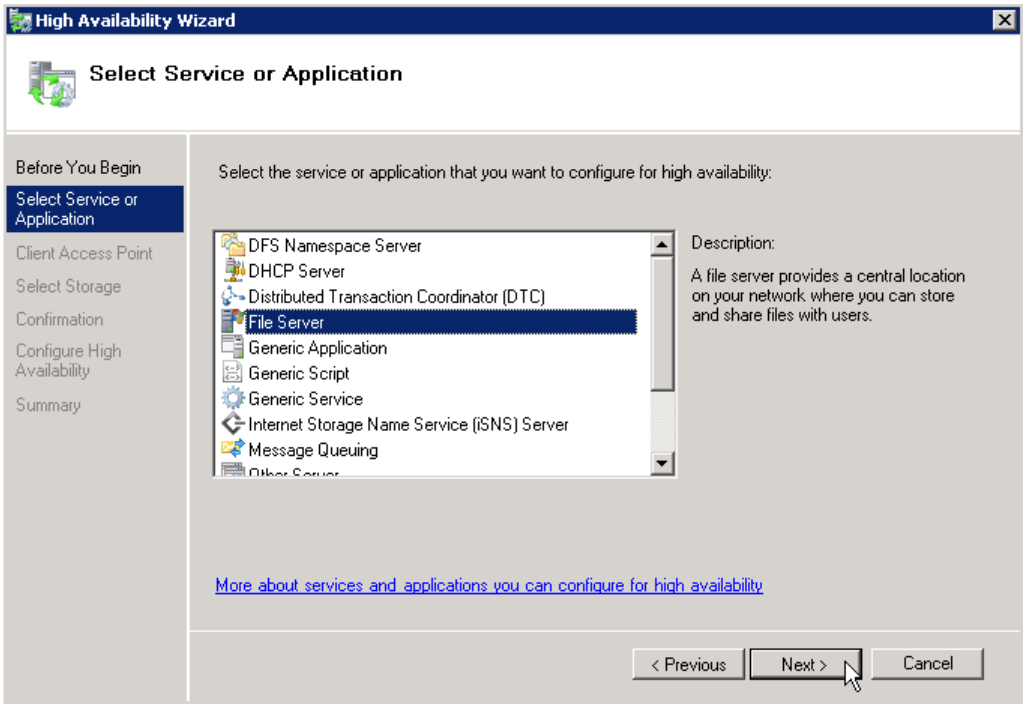

Figure 17 Select Service or Application

5. Follow the instructions in the wizard to specify the following details:
   - A name for the clustered file server
   - The storage volume or volumes that the clustered file server should use

NOTE:

The clustered file server name is the name of the server that users should use to access their file content.

6. On the **Summary** page, review the configuration details and then click **Finish**.
7. In the console tree, make sure **Services and Applications** is expanded, and then select the clustered file server that you just created.
8. Under **Actions**, click **Add a shared folder**

   The **Provision a Shared Folder Wizard** appears. This is the same wizard that you would use to provision a share on a nonclustered file server.
9. Follow the instructions in the wizard to specify the following settings for the shared folder:
   - Path and name
   - NTFS permissions (optional)
   - Advanced settings for the SMB protocol (optional)
   - Whether the NFS protocol will be used for support of UNIX-based clients (optional)
10. After completing the wizard, confirm that the clustered file server comes online.

Verify that the cluster is operational

In order to test cluster functionality, move the clustered file server from one server to the other server. When services or applications are moved, they should fail over to the other node in the cluster.

- Right-click the clustered file server, select **Move this service or application to another node**, and click the available choice of node. When prompted, confirm your choice.

Verify that the status changes in the center pane of the Failover Cluster Management snap-in as the clustered file server instance is moved.

4 Cluster administration

One important feature of HP StorageWorks P4000 G2 Unified NAS Gateways is that they can operate as a single node or as a cluster. This chapter discusses cluster installation and management issues.

Cluster overview

Up to eight server nodes can be connected to each other and deployed as a no single point of failure (NSPOF) cluster. Utilizing a private network allows for communication amongst servers, allowing you to track the state of each cluster node. Each node sends out periodic messages to the other nodes; these messages are called heartbeats. If a node stops sending heartbeats, the cluster service fails over any resources that the node owns to another node. For example, if the node that owns the Quorum disk is shut down for any reason, its heartbeat stops. The other nodes detect the lack of the heartbeat and another node takes over ownership of the Quorum disk and the cluster.

Clustering servers greatly enhances the availability of file serving by enabling file shares to fail over to an alternative server if problems arise. Clients see only a brief interruption of service as the file share resource transitions from one server node to the other.

Cluster terms and components

Nodes

The most basic parts of a cluster are the servers, referred to as nodes. A server node is any individual server in a cluster, or a member of the cluster.

Resources

Hardware and software components that are managed by the cluster service are called cluster resources. Cluster resources have three defining characteristics:

- They can be brought online and taken offline.
- They can be managed in a cluster.
- They can be owned by only one node at a time.

Some resources are created automatically by the system and other resources must be set up manually. Resource types include:

- IP address resource
- Cluster name resource
- Cluster quorum disk resource
- Physical disk resource
- Virtual server name resources
- CIFS file share resources
- NFS file share resources

- FTP file share resources
- iSCSI resources

Cluster groups

Cluster resources are placed together in cluster groups. Groups are the basic unit of failover between nodes. Resources do not fail over individually; they fail over with the group in which they are contained.

Virtual servers

A virtual server is a cluster group that consists of a static IP Address resource and a Network Name resource. Several virtual servers can be created. By assigning ownership of the virtual servers to the different server nodes, the processing load on the storage systems can be distributed between the nodes of a cluster.

The creation of a virtual server allows resources dependent on the virtual server to fail over and fail back between the cluster nodes. Cluster resources are assigned to the virtual server to ensure non-disruptive service of the resources to the clients.

Failover and failback

Failover of cluster groups and resources happens:

- When a node hosting the group becomes inactive.
- When all of the resources within the group are dependent on one resource, and that resource fails.
- When an administrator forces a failover.

A resource and all of its dependencies must be located in the same group so that if a resource fails over, all of its dependent resources fail over.

When a resource is failed over, the cluster service performs certain procedures. First, all of the resources are taken offline in an order defined by the resource dependencies. Secondly, the cluster service attempts to transfer the group to the next node on the preferred owner's list. If the transfer is successful, the resources are brought online in accordance with the resource dependency structure.

The system failover policy defines how the cluster detects and responds to the failure of individual resources in the group. After a previously failed node comes online, the cluster service can fail back the groups to the original host, depending on the failback policy setting. . The failback policy must be set before the failover occurs so that failback works as intended.

Quorum disk

Each cluster must have a shared disk called the Quorum disk. The Quorum disk is the shared storage used by the cluster nodes to coordinate the internal cluster state. This physical disk in the common cluster disk array plays a critical role in cluster operations. The Quorum disk offers a means of persistent storage. The disk must provide physical storage that can be accessed by all nodes in the cluster. If a node has control of the quorum resource upon startup, it can initiate the cluster. In addition, if the node can communicate with the node that owns the quorum resource, it can join or remain in the cluster.

The Quorum disk maintains data integrity by:

- Storing the most current version of the cluster database.
- Guaranteeing that only one set of active communicating nodes is allowed to operate as a cluster.

Cluster planning

Successful cluster planning includes:

- Storage planning
- Network planning
- Protocol planning

Storage planning

For clustering, a basic disk must be designated for the cluster and configured as the Quorum disk.

Additional basic disks are presented to each cluster node for data storage as physical disk resources. The physical disk resources are required for the basic disks to successfully work in a cluster environment, protecting it from simultaneous access from each node.

The basic disk must be added as a physical disk resource to an existing cluster group or a new cluster group needs to be created for the resource. Cluster groups can contain more than one physical disk resource depending on the site-specific requirements.

NOTE:

The LUN underlying the basic disk should be presented to only one node of the cluster using selective storage presentation or SAN zoning, or having only one node online at all times until the physical resource for the basic disk is established.

In preparing for the cluster installation:

- All shared disks, including the Quorum disk, must be accessible from all nodes. When testing connectivity between the nodes and the LUN, only one node should be given access to the LUN at a time.
- All shared disks must be configured as basic (not dynamic).
- All partitions on the disks must be formatted as NTFS.

Network planning

Clusters require more sophisticated networking arrangements than stand alone storage systems. A Windows NT domain or Active Directory domain must be in place to contain the cluster names, virtual server names, and user and group information. A cluster cannot be deployed into a non-domain environment.

All cluster deployments have at least six network addresses and four network names:

- The cluster name (Unique NETBIOS Name) and IP address
- Node A's name and IP address
- Node B's name and IP address
- At least one virtual server name and IP address for virtual server A
- Cluster Interconnect static IP addresses for Node A and Node B

In multi-node deployments, additional network addresses are required. For each additional node, three static IP addresses are required.

Virtual names and addresses are the only identification used by clients on the network. Because the names and addresses are virtual, their ownership can transition from one node to the other during a failover, preserving access to the resources in the cluster group.

A cluster uses at least two network connections on each node:

- The private cluster interconnect or "heartbeat" crossover cable connects to one of the network ports on each cluster node. In more than two node deployments, a private VLAN on a switch or hub is required for the cluster interconnect.
- The public client network subnet connects to the remaining network ports on each cluster node. The cluster node names and virtual server names have IP addresses residing on these subnets.

NOTE:

If the share is to remain available during a failover, each cluster node must be connected to the same network subnet. It is impossible for a cluster node to serve the data to a network to which it is not connected.

Protocol planning

Not all file sharing protocols can take advantage of clustering. If a protocol does not support clustering, it will not have a cluster resource and will not failover with any cluster group. In the case of a failover, a client cannot use the virtual name or virtual IP address to access the share since the protocol cannot failover with the cluster group. The client must wait until the initial node is brought back online to access the share.

HP recommends placing cluster aware and non cluster aware protocols on different file shares.

Table 4 Sharing protocol cluster support

Protocol	Client Variant	Cluster Aware (supports failover)	Supported on cluster nodes
CIFS/SMB	Windows NT Windows 2000 Windows 95 Windows 98 Windows ME	Yes	Yes
NFS	UNIX Linux	Yes	Yes
HTTP	Web	No	Yes
FTP	Many	Yes	Yes
NCP	Novell	No	Yes
AppleTalk	Apple	No	No
iSCSI	Standards-based iSCSI initiator	Yes	Yes

NOTE:

AppleTalk is not supported on clustered disk resources. AppleTalk requires local memory for volume indexing. On failover events, the memory map is lost and data corruption can occur.

Cluster groups and resources, including file shares

The Failover Cluster Management tool (**Start > Administrative Tools > Failover Cluster Management)** provides complete online help for all cluster administration activities.

Cluster resources include administrative types of resources as well as file shares. The following sections include overview and planning issues for cluster groups, cluster resources, and clustered file shares.

Creating and managing these resources and groups must be managed through Failover Cluster Management.

Cluster group overview

A default cluster group is automatically created when the cluster is first created. This default cluster group contains an Internet Protocol (IP) Address resource, a Network Name resource, and the Quorum disk resource. When the new cluster is created, the (IP) address and the cluster name that were specified during setup are set up as the IP address and network name of this default cluster group.

△ CAUTION:

Do not delete or rename the Cluster Group or IP Address. Doing so results in losing the cluster and requires reinstallation of the cluster.

When creating groups, the administrator's first priority is to gain an understanding of how to manage the groups and their resources. Administrators may choose to create a resource group and a virtual server for each node that will contain all resources owned by that node, or the administrator may choose to create a resource group and virtual server for each physical disk resource. Additionally, the administrator should try to balance the load of the groups and their resources on the cluster between the nodes.

Node-based cluster groups

Creating one resource group and one virtual server for each node facilitates group and resource administration. This setup allows administrators to include all file share resources under one group. Clients access all of the resources owned by one node through a virtual server name.

In node-based cluster groups, each group has its own network name and IP address. The administrator decides on which node to place each physical disk resource. This configuration provides a very coarse level of granularity. All resources within a group must remain on the same node. Only two IP addresses and network names are required. This configuration creates less overhead for resource and network administration. A possible disadvantage of this approach is that the resource groups can potentially grow large when many file shares are created.

Load balancing

The creation of separate cluster groups for each virtual server provides more flexibility in balancing the processing load on the cluster between the two nodes. Each cluster group can be assigned to a cluster node using the preferred owner parameter. For example, if there are two cluster groups, the cluster could be set up to have the first cluster group owned by Node A and the second cluster group owned by Node B. This allows the network load to be handled by both devices simultaneously. If only one cluster group exists, it can only be owned by one node and the other node would not serve any network traffic.

File share resource planning issues

CIFS and NFS are cluster-aware protocols that support the Active/Active cluster model, allowing resources to be distributed and processed on both nodes at the same time. For example, some NFS file share resources can be assigned to a group owned by a virtual server for Node A and additional NFS file share resources can be assigned to a group owned by a virtual server for Node B.

Configuring file shares as cluster resources provides for high availability of file shares. Because the resources are placed into groups with a virtual server, ownership of the files can easily move from one node to the other, as circumstances require. If the cluster node owning the group of file shares should be shut down or fail, the other node in the cluster will begin sharing the directories until the original owner node is brought back on line. At that time, ownership of the group and its resources can be brought back to the original owner node.

Resource planning

1. Create a cluster group for each node in the cluster with an IP address resource and a network name resource.

   Cluster resource groups are used to balance the processing load on the servers. Distribute ownership of the groups between the virtual servers.

2. For NFS environments, configure the NFS server.

   NFS-specific procedures include entering audit and file lock information as well as setting up client groups and user name mappings. These procedures are not unique to a clustered deployment and are detailed in the Microsoft Services for NFS section within the “Other network file and print services” chapter. Changes to NFS setup information are automatically replicated to all nodes in a cluster.

3. Create the file share resources.
4. Assign ownership of the file share resources to resource groups.
   a. Divide ownership of the file share resources between the resource groups, which are in turn distributed between the virtual servers, for effective load balancing.
   b. Verify that the physical disk resource for this file share is also included in this group.
   c. Verify that the resources are dependent on the virtual servers and physical disk resources from which the file share was created.

Permissions and access rights on share resources

File share and NFS Share permissions must be managed using the Failover Cluster Management tool and not through the individual shares on the file system themselves via Windows Explorer. Administering them through Failover Cluster Management allows permissions to migrate from one node to other. In contrast, permissions established using Explorer are lost if the share is failed or taken offline.

NFS cluster-specific issues

- Back up user and group mappings.

  To avoid loss of complex advanced mappings in the case of a system failure, back up the mappings whenever the mappings have been edited or new mappings have been added.

- Map consistently.

  Groups that are mapped to each other should contain the same users and the members of the groups should be properly mapped to each other to ensure proper file access.

- Map properly.
  - Valid UNIX users should be mapped to valid Windows users.
  - Valid UNIX groups should be mapped to valid Windows groups.
  - Mapped Windows users must have the "Access this computer from the Network privilege" or the mapping will not function properly.
  - Mapped Windows users must have an active password, or the mapping will not function properly.

- In a clustered deployment, create user name mappings using domain user accounts.

  Because the security identifiers of local accounts are recognized only by the local server, other nodes in the cluster will not be able to resolve those accounts during a failover. Do not create mappings using local user and group accounts.

- In a clustered deployment, administer user name mapping on a computer that belongs to a trusted domain.

  If NFS administration tasks are performed on a computer that belongs to a domain that is not trusted by the domain of the cluster, the changes are not properly replicated among the nodes in the cluster.

- In a clustered deployment, if PCNFS password and group files are being used to provide user and group information, these files must be located on each node of the system.

  Example: If the password and group files are located at `c:\maps` on node 1, then they must also be at `c:\maps` on node 2. The contents of the password and group files must be the same on both nodes as well.

  These password and group files on each server node must be updated periodically to maintain consistency and prevent users or groups from being inadvertently squashed.

Non-cluster aware file sharing protocols

Services for Macintosh (SFM), File and Print Services for NetWare, HTTP file sharing protocols are not cluster aware and will experience service interruption if installed on a clustered resource during failover events of the resource. Service interruptions will be similar to those experienced during a server outage. Data that has not been saved to disk prior to the outage will experience data loss. In the case of SFM, it is not supported because SFM maintains state information in memory. Specifically, the Macintosh volume index is located in paged pool memory. Using SFM in clustered mode is not supported and may result in data loss similar in nature to a downed server should the resource it is based on fails over to the opposing node.

Adding new storage to a cluster

Present the new storage to one node in the cluster. This can be accomplished through selective storage presentation or through SAN zoning.

The tasks described below are used to add storage to a cluster. See the online help for clustering for additional details.

Creating physical disk resources

A physical disk resource must reside within a cluster group. An existing cluster group can be used or a new cluster group may be created. For information on creating disk resources, see the cluster online help topic *Physical Disk resource type*.

NOTE:

- Physical disk resources usually do not have any dependencies set.
- In multi-node clusters it is necessary to specify the node to move the group to. When a cluster group is moved to another node, all resources in that group are moved.
- When a physical disk resource is owned by a node, the disk appears as an unknown, unreadable disk to all other cluster nodes. This is a normal condition. When the physical disk resource moves to another node, the disk resource then becomes readable.

Creating file share resources

To create a file share resource, see two clustering online help topics:

- Create a cluster-managed file share
- Using a server cluster with large numbers of file shares

NOTE:

- A file share resource must reside in the same cluster group as the physical disk resource it will reside on.
- The physical disk resource specified in this step must reside in the same cluster group as specified in the beginning of this wizard.

Creating NFS share resources

To create an NFS share resource, see "MSNFS administration on a server cluster" on page 47.

Shadow copies in a cluster

HP recommends that the cache file be placed on a separate disk from the original data. In this case, a physical disk resource for the cache file disk should be created in the same cluster group as the intended Shadow Copy resource and the volume for which snapshots will be enabled. The resource should be created prior to the establishment of Shadow Copies. The Shadow Copy resource should be dependent on both the original physical disk resource and the physical disk resource that contains the cache file.

For more information, see the following topics in the clustering online help:

- Using Shadow Copies of Shared Folders in a server cluster
- Enable Shadow Copies for shared folders in a cluster

Extend a LUN in a cluster

To extend a LUN on a storage array in a cluster, review the requirements and procedures from the storage array hardware provider for expanding or extending storage.

For additional information associated with extending a LUN in a cluster, see the P4000 SAN documentation at http://www.hp.com/go/p4000.

MSNFS administration on a server cluster

The Microsoft Services for Network File System (NFS) online help provides server cluster information for the following topics:

- Configuring shared folders on a server cluster
  - Configuring an NFS share as a cluster resource
  - Modifying an NFS shared cluster resource
  - Deleting an NFS shared cluster resource
- Using Microsoft Services for NFS with server clusters
  - Understanding how Server for NFS works with server clusters
  - Using Server for NFS on a server cluster
- Configuring User Name Mapping on a server cluster

For further details, see the online help for Microsoft Services for Network File System.

Best practices for running Server for NFS in a server cluster

- Stop Server for NFS before stopping the server cluster.
- Ensure share availability when a node fails.
- Use the appropriate tool to manage Network File System (NFS) share cluster resources.
- Avoid conflicting share names.
- Ensure the availability of audit logs.
- Move file shares or take them offline before stopping Server for NFS.
- Take resources offline before modifying.
- Administer Server for NFS only from computers in a trusted domain.
- Restart the Server for NFS service after the cluster service restarts.
- Choose the appropriate sharing mode.
- Use the command line properly when creating or modifying NFS share cluster resources.
- Use hard mounts.
- Use the correct virtual server name.

Print services in a cluster

The Windows Storage Server 2008 R2 Cluster service implementation increases availability of critical print servers. A print spooler service on a clustered print server may be hosted on any of the nodes in the cluster. As with all cluster resources, clients should access the print server by its virtual network name or virtual IP address.

Creating a cluster printer spooler

Printer spoolers should be created in a separate group dedicated to this purpose for ease of management. For each printer spooler, a physical resource is required to instantiate the print spooler resource. In some cases, dedicated physical resources are not available and hence sharing of the physical resource among other members of the group is acceptable, remembering that all members of a group are managed as a unit. Hence, the group will failover and failback as a group.

To create a printer spooler:

1. Create a dedicated group (if desired).
2. Create a physical resource (disk) (if required, see note).
3. Create an IP address resource for the Virtual Server to be created (if required, see note).
4. Create a Virtual Server Resource (Network Name) (if required, see note).

NOTE:

If the printer spool resource is added to an existing group with a physical resource, IP address, and virtual server resource, steps 1-4 are not required.

5. Create a Print Spool resource.
6. To add a printer to the virtual server:
   a. Double-click the printers and faxes icon.
   b. Right-click the new screen, and then click **add printer**. A wizard starts.
   c. Click **create a new port**, and then click **Next**.
   d. Enter the IP address of the network printer.
   e. Update the Port Name if desired, click **Next**, and then click **Finish**.
   f. Select the appropriate driver, and then click **Next**.
   g. If presented with a dialog to replace the driver present, click **keep the driver**, and then click **Next**.
   h. Name the printer, and then click **Next**.
   i. Provide a share name for the printer for network access, and then click **Next**.
   j. Provide location information and comments, and then click **Next**.
   k. Click **Yes** to print a test page, click **Next**, and then click **Finish**.
   l. A dialog box appears regarding the test page. Select the appropriate answer.

The Printer Spool is now a clustered resource.

Advanced cluster administration procedures

Failing over and failing back

As previously mentioned, when a node goes offline, all resources dependent on that node are automatically failed over to another node. Processing continues, but in a reduced manner, because all operations must be processed on the remaining node(s). In clusters containing more than two nodes, additional fail over rules can be applied. For instance, groups can be configured to fail over different nodes to balance the additional work load imposed by the failed node. Nodes can be

excluded from the possible owners list to prevent a resource from coming online on a particular node. Lastly the preferred owners list can be ordered, to provide an ordered list of failover nodes. Using these tools, the failover of resources can be controlled with in a multinode cluster to provide a controlled balanced failover methodology that balances the increased work load.

Because operating environments differ, the administrator must indicate whether the system will automatically fail the resources (organized by resource groups) back to their original node or will leave the resources failed over, waiting for the resources to be moved back manually.

NOTE:

If the storage system is not set to automatically fail back the resources to their designated owner, the resources must be moved back manually each time a failover occurs.

Restarting one cluster node

CAUTION:

Restarting a cluster node should be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being restarted. Attached connections can be viewed through Server Manager on the storage system Desktop using Terminal Services. From Server Manager, select **File Sharing > Shared Folders > Sessions**.

The physical process of restarting one of the nodes of a cluster is the same as restarting a storage system in single node environment. However, additional caution is needed.

Restarting a cluster node causes all cluster resources served by that node to fail over to the other nodes in the cluster based on the failover policy in place. Until the failover process completes, any currently executing read and write operations will fail. Other node(s) in the cluster will be placed under a heavier load by the extra work until the restarted node comes up and the resources are moved back.

Shutting down one cluster node

CAUTION:

Shutting down a cluster node must be done only after confirming that the other node(s) in the cluster are functioning normally. Adequate warning should be given to users connected to resources of the node being shutdown.

Shutting down a cluster node causes all cluster resources served by that node to fail over to the other node(s). This causes any currently executing client read and write operations to fail until the cluster failover process completes. The other node(s) are placed under a heavier load by the extra work until the second node is powered up and rejoins the cluster.

Powering down the cluster

The power down process for the storage system cluster is similar to the process for a single node, but with the cluster, extra care must be taken with the storage subsystem and the sequence of the shutdown.

The power down process is divided into two main steps:

1. Shutting down the cluster nodes
2. Removing power from the cluster nodes

The sequence of these steps is critical. The devices must be shut down before the storage subsystem. Improperly shutting down the nodes and the storage subsystem causes corruption and loss of data.

△ CAUTION:

Before powering down the cluster nodes, follow the proper shutdown procedure as previously illustrated. See "Shutting down one cluster node." Only one cluster node should be shut down at a time.

Powering up the cluster

The power up process for the storage system cluster is more complex than it is for a single node because extra care must be taken with the storage subsystem.

The sequence of the power up steps is critical. Improper power up procedures can cause corruption and loss of data.

△ CAUTION:

Do not power up the cluster nodes without first powering up the storage subsystem, and verifying it is operating normally.

Nodes should be powered up separately allowing one node to form the cluster prior to powering up the additional node(s). To power up the cluster nodes:

1. After the storage subsystem is confirmed to be operating normally, power up a single node. Wait for the node to come completely up before powering up the subsequent node(s).

   If more than one node is powered up at the same time, the first node that completes the sequence gains ownership of the cluster quorum and controls the cluster database. Designate a particular node as the usual cluster quorum owner by always powering up that node first and letting it completely restart before powering up additional cluster node(s).
2. Power up the additional cluster node(s). Each node should be allowed to start fully, prior to starting a subsequent node.

5 Administration tools

HP StorageWorks P4000 G2 Unified NAS Gateways include several administration tools to simplify storage system management tasks.

Microsoft Windows Storage Server 2008 R2 administration tools

Microsoft® Windows® Storage Server 2008 R2 operating systems provide a user interface for initial server configuration, unified storage system management, simplified setup and management of storage and shared folders, and support for Microsoft iSCSI Software Target. It is specially tuned to provide optimal performance for network-attached storage. Windows Storage Server 2008 R2 provides significant enhancements in share and storage management scenarios, as well as integration of storage system management components and functionality.

Remote Desktop for Administration

You can remotely administer storage systems by using Remote Desktop for Administration (formerly known as Terminal Services in Remote Administration mode). You can use it to administer a computer from virtually any computer on your network. Based on Terminal Services technology, Remote Desktop for Administration is specifically designed for server management.

Remote Desktop for Administration does not require the purchase of special licenses for client computers that access the server. It is not necessary to install Terminal Server Licensing when using Remote Desktop for Administration.

You can use Remote Desktop for Administration to log on to the server remotely with any of the following features:

- Remote Desktop Connection
- Remote Web Administration
- Windows Server Remote Administration Applet

For more information, see the Windows Storage Server 2008 R2 Help.

Share and Storage Management

With the Share and Storage Management snap-in provided in this release, you can more easily set up and manage shared folders and storage. Share and Storage Management provides the following:

- MMC-based management of shared folders and storage.
- Provision Storage Wizard for creating and configuring storage for file sharing and block sharing, including creating LUNs on storage subsystems, as well as creating and formatting volumes on LUNs or server disks.

NOTE:

You must have a VDS Hardware Provider that is appropriate for your storage system installed in order to provision storage on an iSCSI target. If you have Microsoft iSCSI Software Target running on a Windows Storage Server 2008 R2 storage system, install the Microsoft iSCSI Software Target VDS Hardware Provider on the client computer.

- Provision a Shared Folder Wizard for creating and configuring shared folders that can be accessed by using either the server message block (SMB) or NFS protocol.
- Single Instance Storage (SIS) can be enabled or disabled for each volume that is displayed in Share and Storage Management. SIS recovers disk space by reducing the amount of redundant data stored on a volume. It identifies identical files, storing only a single copy of the file in the SIS Common Store, and replacing the files with pointers to the file in the SIS Common Store.

The Share and Storage Management snap-in makes it possible to complete most of the administrative tasks that are required to create and manage shared folders and volumes without having to use the Shared Folder Management, Storage Manager for SANs, or Disk Management snap-ins. These tasks include configuring quotas to restrict the quantity of data, configuring file screening to prevent certain file types or only allowing certain file types defined by the administrator, and enabling indexing.

For more information, see the Windows Storage Server 2008 R2 Help.

Single Instance Storage

The Single Instance Storage (SIS) feature reduces the amount of space that is used to store data on a volume. SIS does this by replacing duplicate files with logical links that point to a single copy of the file in the SIS Common Store, which is a hidden folder that is located in the root directory of the volume.

SIS consists of two primary components that together maintain a database of file signatures. These components include:

- Groveler service - The Groveler service scans the hard-disk volumes on a server for duplicate copies of files. If the service locates duplicate copies of files, the information about the duplicates is sent to the Single Instance Storage Filter. The Groveler service runs as a user-level service.
- Single Instance Storage Filter - The Single Instance Storage Filter is a file system filter service that manages duplicate copies of files on hard-disk volumes. When notified by the Groveler service of duplicate copies of files, this component copies one instance of a duplicate file into a central folder. The duplicate is then replaced by a link (a reparse point) to the central copy. The link file contains information about the original file, such as its current location, size, and attributes. The Single Instance Storage Filter runs in kernel mode.

The Single Instance Storage Filter service cannot be stopped. If this service is disabled, the linked files are not accessible. If the central folder is deleted, the linked files can become permanently inaccessible. If you stop the Groveler service, the files cannot be automatically linked, but the existing linked files can still be accessible.

You can enable SIS on a maximum of 20 volumes per computer. SIS cannot act upon any files that are referenced through junction points, and it cannot be used with any file system except the NTFS file system. SIS will not process files that are 32 kilobytes or less in size.

If you need to access data that is stored on a SIS volume, which might be required for backup and recovery operations, you must either run or have installed Single Instance Storage Filter on your computer.

Backup and recovery by using SIS has the following requirements:

- The backup software used must support SIS-enabled volumes.
- The SIS volume, SIS Common Store folder, and reparse points (links) to the files must be restored to a Windows 2000 NTFS version 5.0 (or later) file system or partition that supports reparse points or junction points.
- The Single Instance Storage Filter must be installed or enabled to access the data in the SIS volume.
- The backup program must be capable and configured to backup and restore the reparse points or junction points (links) to the files, and the SIS volume and the SIS Common Store folder must be selected.

For more information, see the Windows Storage Server 2008 R2 Help.

Print Management

Print Management is an MMC snap-in that you can use to view and manage printers and print servers in your organization. You can use Print Management from any computer running Windows Storage Server 2008 R2, and you can manage all network printers on print servers running Windows 2000 Server, Windows Server 2003, Windows Storage Server 2003, Windows Storage Server 2003 R2, or Windows Storage Server 2008 R2.

Print Management provides details such as the queue status, printer name, driver name, and server name. You can also set custom views by using the Print Management filtering capability. For example, you can create a view that displays only printers in a particular error state. You can also configure Print Management to send e-mail notifications or run scripts when a printer or print server needs attention. The filtering capability also allows you to bulk edit print jobs, such as canceling all print jobs at once. You can also delete multiple printers at the same time.

Administrators can install printers remotely by using the automatic detection feature, which finds and installs printers on the local subnet to the local print server. Administrators can log on remotely to a server at a branch location, and then install printers remotely.

For more information, see the Windows Storage Server 2008 R2 Help.

Network File System (NFS) User Mapping

Network File System (NFS) is a network file sharing protocol that allows remote access to files over a network and is typically used in networks with computers running UNIX, Linux, or Mac OS operating systems. NFS is supported on all HP P4000 G2 Unified NAS Gateways.

All of the following types of NFS account mapping are supported: Active Directory® Domain Services (AD DS) mapped user access, Active Directory® Lightweight Directory Services (AD LDS) mapped user access, unmapped anonymous user access, and unmapped UNIX user access.

For more detailed information about each of these access methods, see the *NFS Account Mapping in Windows Server 2008 R2* whitepaper, which is available for download at:

http://www.microsoft.com/downloads/en/details.aspx?FamilyID=5f4c294c-8692-4235-8236-8ea809ae71f7 .

Configuring AD LDS

To simplify AD LDS mapped user access, HP provides the scripts that Microsoft describes in the *NFS Account Mapping in Windows Server 2008 R2* whitepaper. The whitepaper describes two phases of scripts. Phase 1 scripts are used to install and prepare your system for NFS account mapping. Phase 2 scripts are used to configure specific users and groups for NFS account mapping.

For more detailed information about the Phase 1 and Phase 2 scripts, see the Microsoft Knowledge Base article *Description of scripts to use to simplify user account mapping between a UNIX client and a Windows-based server* at http://support.microsoft.com/kb/973840.

Microsoft hotfix 2222746

HP supports Microsoft hotfix 2222746. This hotfix applies to Windows Storage Server 2008 R2 as well as the products listed in the KB article. The hotfix addresses the following problem: *File permissions are incorrectly set when you share a folder on a Windows Server 2008-based or Windows Storage Server 2008-based NFS server*. For full details of this hotfix, see http://support.microsoft.com/kb/2222746.

The incorrect behavior that the hotfix addresses occurs when using AD LDS mapped user access. HP has installed this hotfix on the storage system but has not enabled it by default. You can enable it by setting the following registry key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\NfsServer\NlmNsm\AutoCorrectPrimaryGroup

Type: REG_DWORD

Value: 0x1

After enabling the hotfix by setting the registry subkey, you must restart the storage system in order for the hotfix to take effect.

Phase 1 scripts

Phase 1 scripts are located in the `c:\hpnas\components\postinstaller\adlds` folder. You enable AD LDS NFS mapping by running **factory-setup-adlds.cmd**. The command is run without any arguments. The script installs two Roles and one Instance:

- Active Directory Lightweight Directory Services (AD LDS) Role
- An AD LDS Instance named **NFSInstance**.
- Services for Network File System (NFS) under the File Services Role

To verify the installation of the AD LDS Role and Instance, in **Administrative Tools**, select **Server Manager**. If the AD LDS Role was successfully installed, you will see **Active Directory Lightweight Directory Services** listed under **Roles**. Click on it and you will see **NFSInstance** listed under **SystemServices** if that instance was successfully installed.

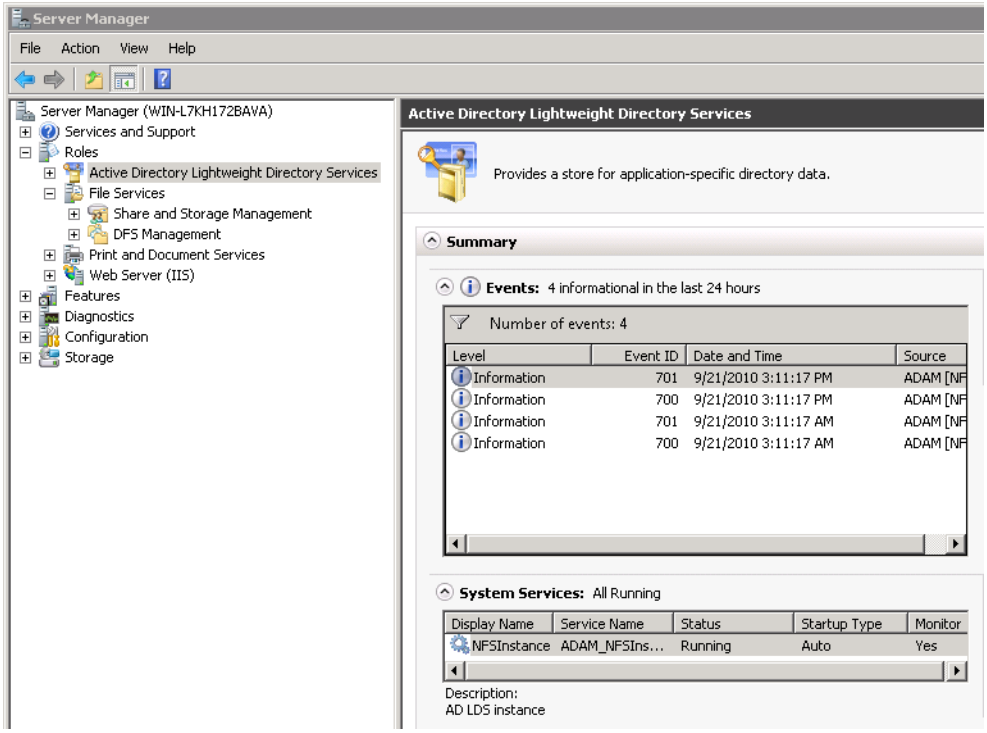

Figure 18 AD LDS Role and Instance

To verify the installation of the Services for Network File System (NFS) Role, in Server Manager, under **Roles**, click **File Services**. In the **System Services** table, **Server for NFS** is listed.

Phase 2 scripts

Phase 2 scripts are located in the `c:\hpnas\components\ADLDS` folder. You create NFS mappings for users and groups by running the **nfs-adlds-config.js** script. This script takes as inputs a standard UNIX password file and a standard UNIX group file. For each user or group in these files a Windows user or group account is created if one does not already exist. Next, the script examines the UNIX password file to extract the User ID and Group ID of each user. From these, a mapping is created on the Windows system that associates the User ID and Group ID with the UNIX account that has the same User ID and Group ID. Similarly, the script examines the UNIX group file, extracts the Group ID, and forms a mapping on the Windows system that associates the Group ID with the UNIX group having the same ID. This mapping is what allows directories and files to be accessed from either the Windows NFS server or the UNIX NFS client using the same User ID and Group ID.

The **nfs-adlds-config.js** script will also add Windows users to the appropriate Windows groups for the newly created user and group accounts. It does this by examining the relationships between users and groups in the password and group files that were given as inputs to the script.

Password and Group file syntax

You can create the password and group files yourself or copy them from the NFS client at **/etc/passwd** and **/etc/group**.

Each line of a standard UNIX password file follows this format:

```
user:password:UID:GID:comment:home directory:command shell
```

All fields are required, but the only fields that are used are the `user`, `UID`, and `GID` fields. If you are creating these files yourself, you may want to leave the other fields blank.

Each line of a standard UNIX group file follows this format:

```
group:password:GID:group list
```

All fields are required, but only the `group`, `GID`, and `group list` fields are used. The `GID` field value must match the `GID` field value in the password file for those users that belong to the group.

NOTE:

If you create the group and password files, you must have corresponding users and groups on the UNIX system. The correspondence is through the numeric UID and GID; however, the user names and group names can be different. For example, the UNIX root group can be associated with a Windows group named **rootgroup** as long as its group ID of 0 is the same between them.

IMPORTANT:

- User names in the password file cannot match group names in the group file. Windows does not allow user names and group names to be the same. An example of this is the root user which typically belongs to the root group on a UNIX system. You would need to rename one of these. For example, in the group file, you might rename the **root** group to **rootgroup**.
- User and group names in Windows are case insensitive. If the password or group files contain accounts whose names differ only in their case, you will need to delete or rename entries in those files.
- Users within the password file must have unique user IDs. Groups within the group file must have unique group IDs.
- All users included in the password file are imported. Consider editing the file before running the configuration script to retain only the users that you want mapped.
- All groups in the group file are imported. Consider editing the file before running the configuration script to retain only the groups that you want mapped.

Script execution

You configure NFS mapping for AD LDS by executing the **nfs-adlds-config.js** script that is located in the `c:\hpnas\components\ADLDS` folder. Executing the script with no command line options will display the following help screen:

```
nfs-adlds-config.js /passwd:passwdfile /group:groupfile /ldf:out.ldf
          /usercmd:generateusers.cmd [/execute] [/log:logname

/passwd       - location of passwd file
/group        - location of group file
/ldf          - output generated ldf file
/usercmd      - output cmd file to generate local users and groups
/userpassword - provide a password to be used for all user accounts created
/execute      - import user objects to ADAM using the generated files
/log          - specify a filename to log operations

Both /passwd and /group must be specified.
At least one of /ldf or /usercmd must be specified.
If /userpassword is not specified all local accounts created must have
passwords set manually before NFS mapping will succeed.
```

Figure 19 AD LDS script execution help screen

As a best practice, specify all of the above parameters so that Windows accounts and NFS mappings are created; however, you can provide finer control as follows. If you omit the `/ldf` option, the script creates Windows accounts but not NFS mappings. Likewise, omitting the `/usercmd` option creates NFS mappings but not Windows accounts.

The `/execute` option controls whether Windows accounts and NFS mappings are actually made to the system. Omitting the `/execute` option will still create the output files (these are the `generateusers.cmd` and `ldf.out` files as shown above). You can then examine these files to see the actions that would have been performed had the `/execute` option been included.

The `/userpassword` option specifies the password that the script assigns when creating Windows accounts for new users. You must use a password that meets the password strength requirements of your system. By default Windows Storage Server 2008 R2 requires strong passwords. If you specify a password that does not meet the requirements, the script will not inform you. It will create the Windows account with a blank password and establish the NFS mapping. Until you change the password to a non-blank value, NFS mapping is disabled for that Windows account. Any attempts in UNIX to use that mapped user will result in **Permissions Denied** or **Input/Output** errors.

NOTE:

If users are created with a blank password because the `/userpassword` option was not given or the specified password does not meet the password strength requirements of the system, those users can log into the system console without the security that a strong password provides.

Verifying script execution

The following steps describe how to verify that the proper Windows accounts were created and the NFS mappings were made when the **nfs-adlds-config.js** script is run. If there are problems, examine the log file specified by using the `/log` option.

After the script is successfully executed, the users in the password file exist as Windows Users and the groups from the group file exist as Windows Groups. You can verify this with Server Manager:

1. Click **Start**, right click **Computer**, and then select **Manage**.
2. Expand the **Configuration** and **Local Users and Groups** nodes.

The imported users and groups are listed in the **Users** and **Groups** folders, respectively.

The newly-created NFS mappings are stored as Active Directory objects and can be verified as follows:

1. Click **Start > Administrative Tools > ADSI Edit**.
2. On the **Action** menu, select **Connect to...**.
3. In the **Connection Settings** dialog box, under **Computer**, select the radio button labeled **Select or type a domain or server: (Server|Domain[:port])**.
4. In the **Select or type a domain or server: (Server|Domain[:port])** field, type **localhost:389** and then click **OK**.

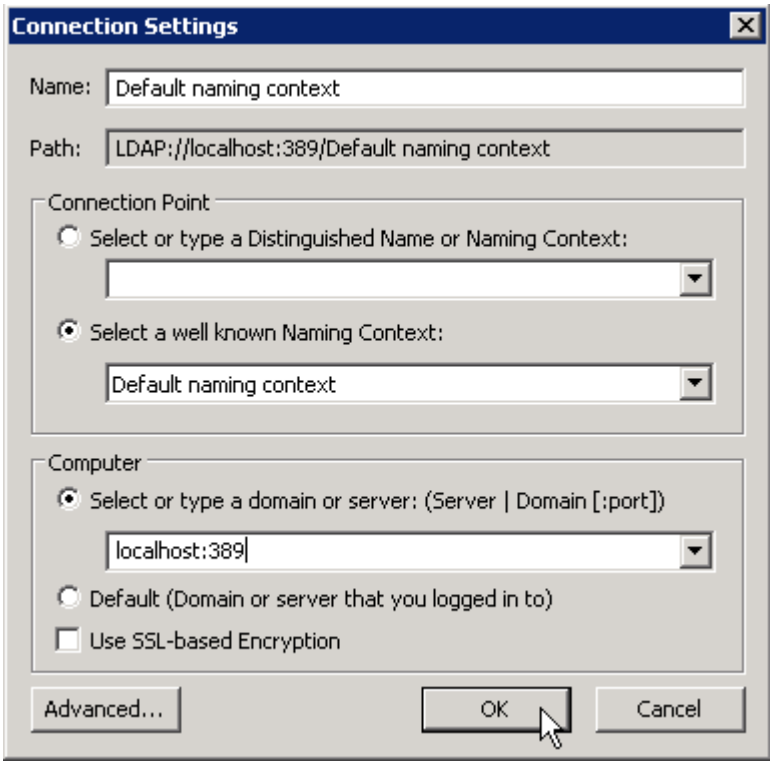

Figure 20 ADSI Edit Connection Settings dialog box

5. Expand the **Default naming context [localhost:389]** node to open **CN=nfs, DC=nfs**, and then **CN=Users**.

The list of NFS-mapped users and groups appears under **CN=Users**. In the figure below, **root**, **john** and **mary** are NFS mapped users. **rootgroup** and **users** are NFS mapped groups. These users and groups must also exist on the UNIX system in order for NFS mapping to work correctly between UNIX and Windows.

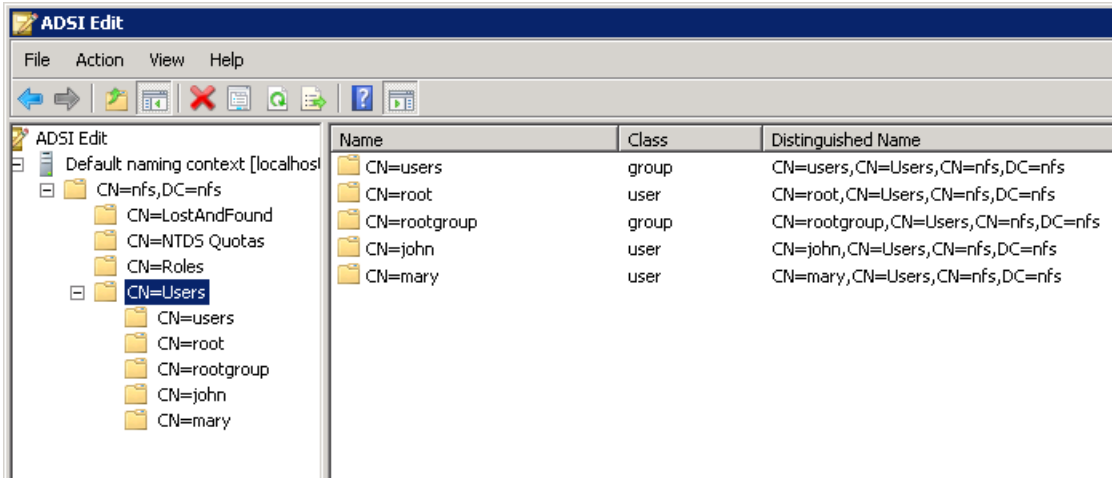

Figure 21 NFS-mapped users and groups in ADSI Edit

Because the imported users and groups are now Windows users and groups as well as UNIX users and groups, you can use NFS sharing so that volumes, folders and files are visible in both the Windows file system and the UNIX file system. When you set ownership or permissions in the Windows file system, the proper ownership and permissions are set on the UNIX side. Likewise, setting ownership or permissions in the UNIX file system results in proper values on the Windows file system.

Shared access example

The following example illustrates how to use the provided AD LDS scripts. By following the procedures, you will create a password and group file that serves as input for the **nfs-adlds-config.js** script. You will then create a Windows folder that you will set to be NFS shared by the group **Everyone**. You will then mount this folder in UNIX and observe how a file created in UNIX is owned by the corresponding mapped user on the Windows system. Similarly you will create a file in Windows, change its ownership to an NFS mapped user, and observe that it is owned by the proper user and group in UNIX. Finally, you will restrict the properties of the Windows shared folder so that it is accessible by a single user instead of the group **Everyone**.

NOTE:

This example assumes that the system is not part of an Active Directory domain or is part of an Active Directory domain but will be using AD LDS for user name mapping.

IMPORTANT:

To ensure proper NFS user name mapping behavior when AD LDS is used for user name mapping, you must enable Microsoft hotfix 2222746, which is installed on the storage system but is not enabled it by default. You can enable it by setting the following registry subkey:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\NfsServer\NlmNsm\AutoCorrectPrimaryGroup

Type: REG_DWORD

Value: 0x1

After setting the registry subkey, restart the storage system.

Step 1

If you have not already run the Phase 1 script **factory-setup-adlds.cmd**, do so now. The script is located in the `c:\hpnas\components\postinstaller\adlds` folder.

Step 2

On the Windows system, in the `C:\hpnas\components\ADLDS` folder, create a file named *passwdfile* containing:

```
root:x:0:0:root:/root:/bin/bash
user1:x:2701:5700:A sample user:/home/user1:/bin/sh
```

Create a file named *groupfile* containing:

```
rootgroup:x:0:
allusers:x:5700:
```

On the UNIX system, edit the **/etc/passwd** file to have an account for user1 with a user ID of 2701 and a group ID of 5700. Edit the **/etc/group** file to have a group account for allusers with a group ID of 5700. On the UNIX side you do not need to create a group named **rootgroup** because the existing group named **root** has a group ID of 0.

NOTE:

If you are taking the **/etc/group** file from the UNIX system and preparing it as input for **nfs-adlds-config.js**, you must rename the root group to something else such as *rootgroup*. This solves the conflict with the root user. As described earlier, Windows does not allow the same name to be used for both a user and a group.

Step 3

Execute the phase 2 script c:\hpnas\components\ADLDS in a command window as follows:

```
nfs-adlds-config.js /passwd:passwdfile /group:groupfile /ldf:out.ldf /
usercmd:generateusers.cmd /userpassword:c0mpleX! /execute /log:log.txt
```

Examine *log.txt* to ensure that there are no errors. Error messages that read **System error 1379: The specified local group already exists** can be ignored.

Follow the steps in "Verifying script execution" on page 57 to verify that appropriate Windows user and group accounts are created and expected NFS mappings are established.

Step 4

Create a folder **C:\NfsTest**. Right-click the folder and select **Properties**. Click the **NFS Sharing** tab and then click the **Manage NFS Sharing** button. Check the **Share this folder** check box. In the **Share name** field, type **NfsTest**. Clear the **Enable unmapped user access** check box.

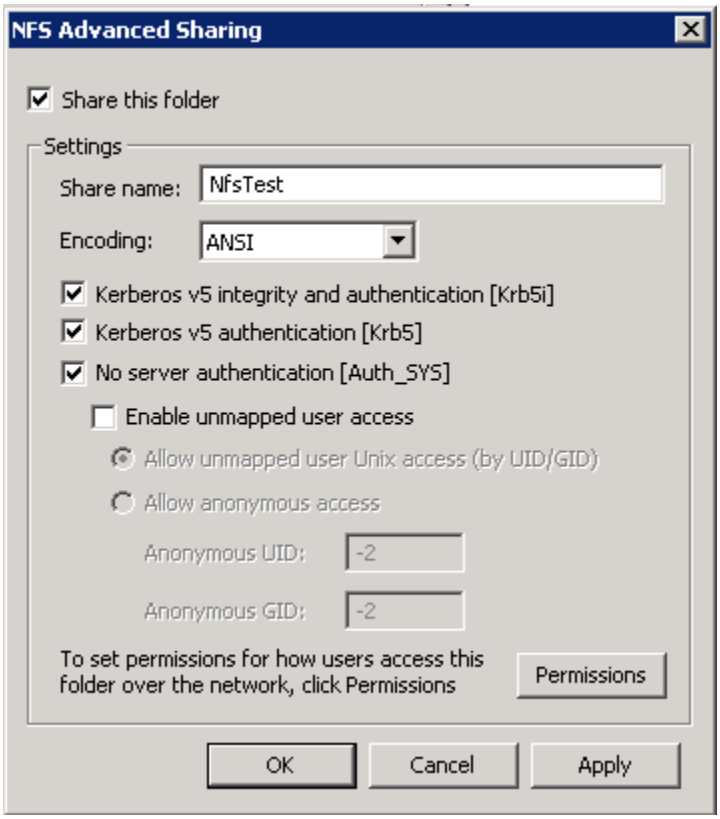

Figure 22 NFS Advanced Sharing dialog box

Click the **Permissions** button. In the **Type of access** list, select **Read-Write**. Check the **Allow root access** check box.

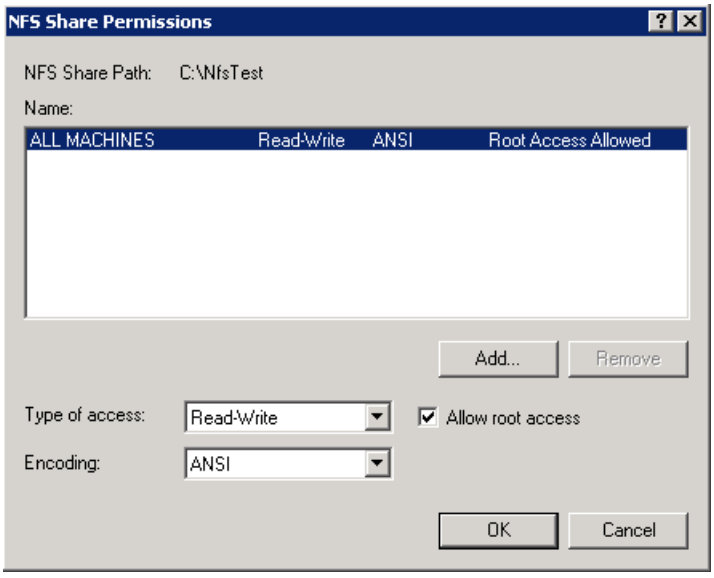

Figure 23 NFS Share Permissions dialog box

Click **OK** twice to return to the NfsTest **Properties** dialog box.

Step 5

In the NfsTest **Properties** dialog box, select the **Security** tab. This tab shows the current security settings for the folder. The following steps will add permissions for **Everyone** to have access to the NfsTest folder. To do this, click **Edit…** and then click **Add…**. Under **Group or user names:** type **Everyone** and then click **OK**. In the **Permissions for NfsTest** dialog box, check the **Allow** check box for **Full control** under **Permissions for Everyone**.

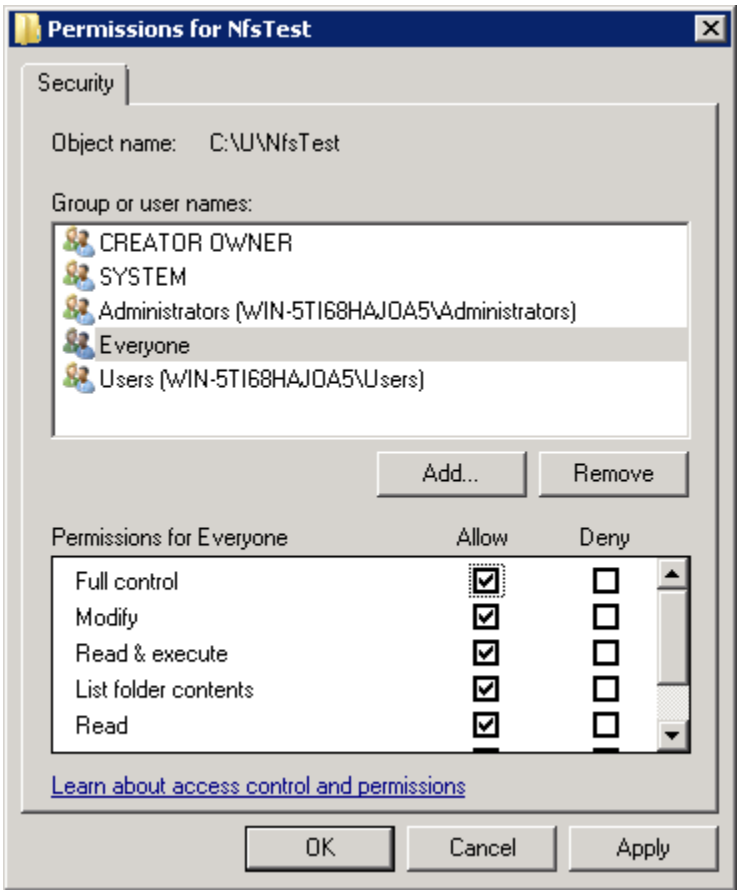

Figure 24 Permissions for NfsTest dialog box

Click **OK** to dismiss the **Permissions for NfsTest** dialog box and then click **OK** to dismiss the **NfsTest Properties** dialog box.

Step 6

On the UNIX system as root, mount the NFS share. A typical command is as follows where 10.30.15.20 is the IP address of the Windows system:

```
> mkdir /mnt/nfstest
```

```
> mount -t nfs 10.30.15.20:/NfsTest /mnt/nfstest
```

As root, create a file in the mounted directory:

```
> touch /mnt/nfstest/rootfile
```

Step 7

On the Windows system, open Windows Explorer to C:\NfsTest. Properties for *rootfile* are displayed. To display the owner of *rootfile*, add the **Owner** column: right-click the column header, select **More...**,

check the **Owner** check box and then click **OK**. Note that *root* is part of the owner name. This verifies that NFS mapping is functional for the root user.

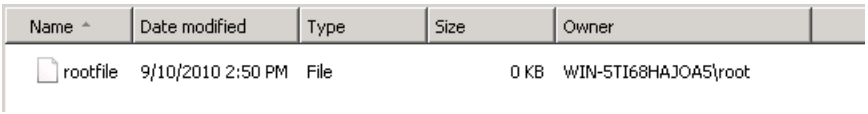

Step 8

On the Windows system as Administrator, create a file file.txt in C:\NfsTest. Windows Explorer now displays its properties:

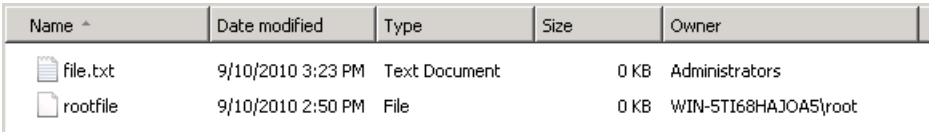

On the UNIX system, issue a listing of the **/mnt/nfstest** directory. This listing displays the following information. Note that file permissions on your system may be different than those shown here.

```
-rwx------ 1 nfsnobody nfsnobody 0 Sep 10 16:23 file.txt
-rw-r--r-- 1 root      root      0 Sep 10 15:50 rootfile
```

The reason that the owner of *file.txt* is **nfsnobody** is that the file is owned by the **Windows Administrators** group, which is not an NFS mapped group.

Step 9

In this step, you will change the owner of *file.txt* to **user1**. On the Windows system, in Windows Explorer, right-click **file.txt**, select **Properties** and then select the **Security** tab. Click **Advanced**, select the **Owner** tab, and then click **Edit...** . Click **Other users or groups...** , type **user1** in the dialog box, and then click **OK**.

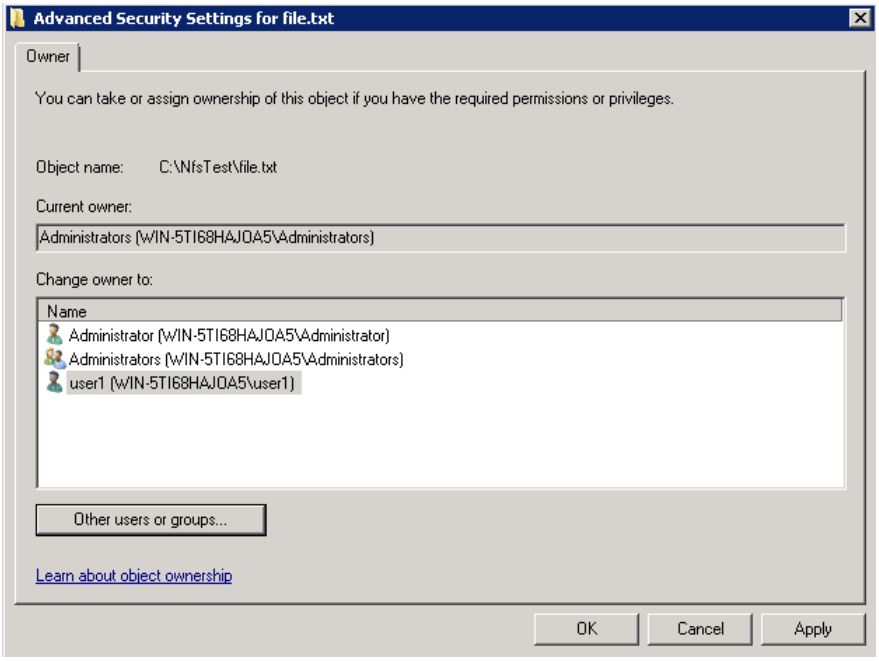

Figure 25 Advanced Security Settings for file.txt

Click **OK** three times to return to Windows Explorer. The file *file.txt* is now owned by **user1**.

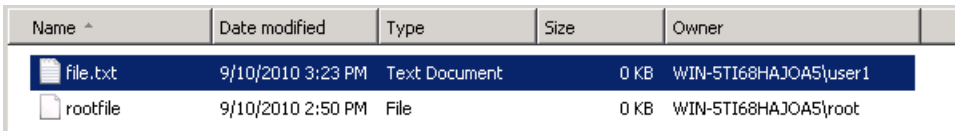

Step 10

On the UNIX system, issue a listing of the **/mnt/nfstest** directory. Because you changed the owner of *file.txt* on the Windows system to **user1**, the owner of the file on the UNIX side is also **user1**. Because **user1** is in the **allusers** group on the Windows system, the group ownership of *file.txt* is **allusers** on the UNIX side.

```
-rw-r--r-- 1 user1 allusers 0 Oct 11 14:18 file.txt
-rw-r--r-- 1 root  root     0 Oct 11 14:17 rootfile
```

If the group for *file.txt* is listed as **nfsnobody**, you must enable Microsoft hotfix 2222746 and restart the storage system as described in "Shared access example" on page 59.

Step 11

In this step, you will modify the permissions on the shared folder to restrict access to the shared folder to a group smaller than **Everyone** by changing the ownership of the shared folder and then specifying permissions for that owner.

First, remove the file *rootfile* from **/mnt/nfstest** on the UNIX system.

In Windows Explorer, right-click the **NfsTest** folder, select **Properties**, and then select the **Security** tab. Click **Advanced**, select the **Owner** tab, and then click **Edit…**. Click **Other users or groups…** and type **user1** in the **Enter the object name to select** field.

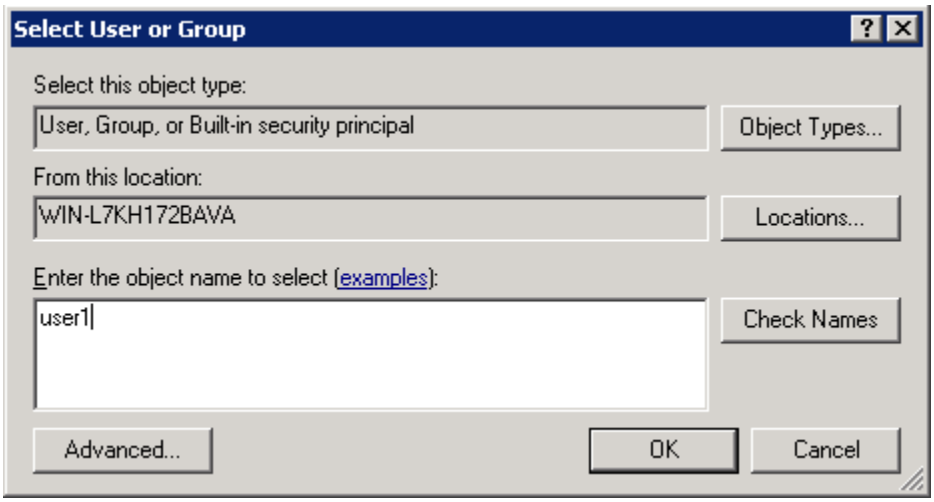

Figure 26 Select User or Group dialog box

Click **OK** to return to the **Advanced Security Settings for NfsTest** dialog box. Check the **Replace owner on subcontainers and objects** check box.

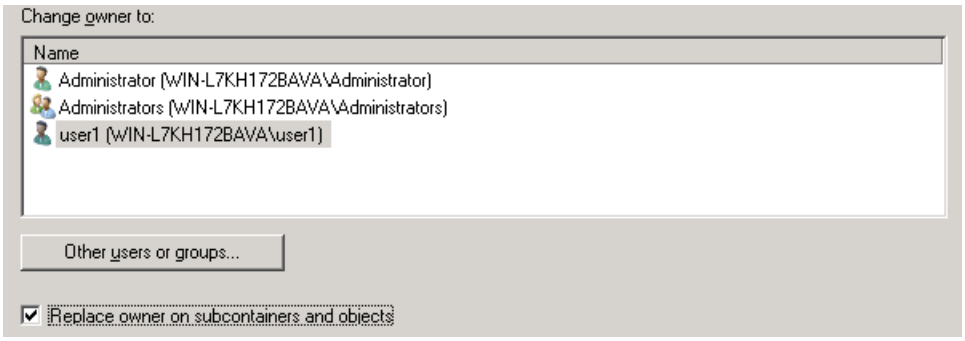

Figure 27 Replace owner on subcontainers and objects

Click **OK** four times to dismiss the **Properties** dialog box. Return to the **Properties** dialog box, select the **Security** tab, and then click **Edit**. Select **Everyone** in the list of groups and user names and then click **Remove**. Next, click **Add** and add **user1**. Select **full control** for its permissions.

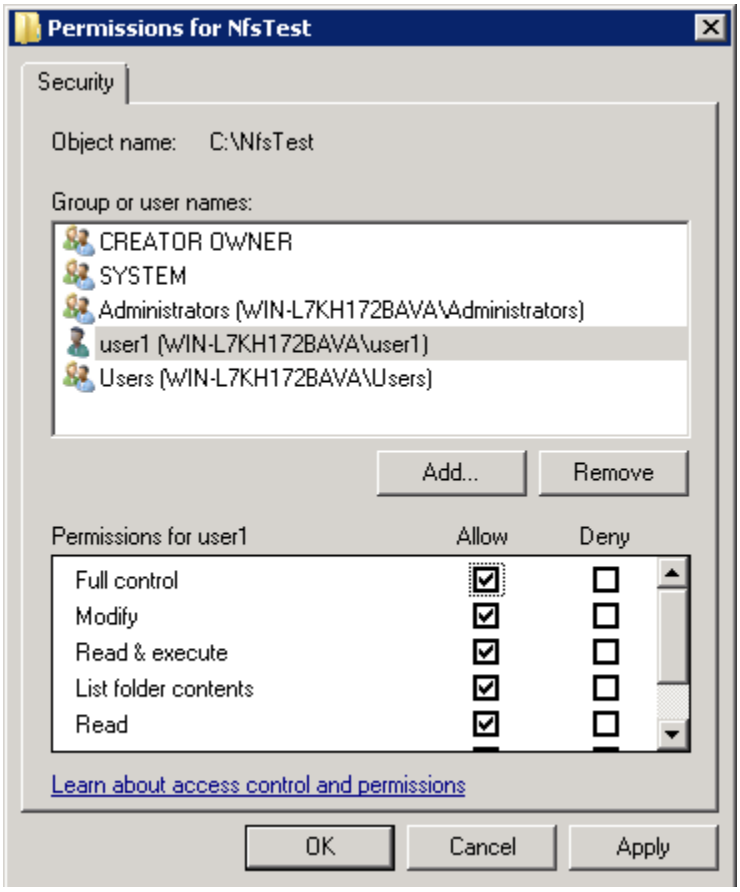

Figure 28 Permissions for NfsTest dialog box

Click **OK** twice to dismiss the NfsTest **Permissions**and **Properties** dialog boxes.

On the UNIX side, you can now issue the `su user1` command, then issue the `cd /mnt/nfstest` command, and create files in that directory.

You can also access the mounted directory as root. If you want to restrict this, return to the **NFS Share Permissions** dialog box as shown in Figure 23 on page 61 and clear the **Allow root access** check box.

6 File server management

This chapter begins by identifying file services in Windows Storage Server 2008 R2. The remainder of the chapter describes the many tasks and utilities that play a role in file server management.

File services features in Windows Storage Server 2008 R2

Single Instance Storage

Single Instance Storage (SIS) provides a copy-on-write link between multiple files. Disk space is recovered by reducing the amount of redundant data stored on a server. If a user has two files sharing disk storage by using SIS, and someone modifies one of the files, users of the other files do not see the changes. The underlying shared disk storage that backs SIS links is maintained by the system and is only deleted if all the SIS links pointing to it are deleted. SIS automatically determines that two or more files have the same content and links them together.

File Server Resource Manager

File Server Resource Manager is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. By using File Server Resource Manager, administrators can place quotas on volumes, actively screen files and folders, and generate comprehensive storage reports.

By using File Server Resource Manager, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and to generate notifications when the quota limits are approached and exceeded.
- Create file screens to screen the files that users can save on volumes and in folders and to send notifications when users attempt to save blocked files.
- Schedule periodic storage reports that allow users to identify trends in disk usage and to monitor attempts to save unauthorized files, or generate the reports on demand.

Windows SharePoint Services

Windows SharePoint Services is an integrated set of collaboration and communication services designed to connect people, information, processes, and systems, within and beyond the organization firewall.

File services management

Information about the storage system in a SAN environment is provided in the HP StorageWorks SAN Manuals page located on the HP web site at www.hp.com/go/SDGManuals.

Configuring data storage

The HP StorageWorks P4000 G2 Unified NAS Gateway is configured only for the operating system. The administrator must configure data storage for the storage system.

Configuring additional data storage involves creating arrays, logical disks, and volumes. Table 5 shows the general task areas to be performed as well as the utilities needed to configure storage for an HP Smart Array-based storage system.

Table 5 Tasks and utilities needed for storage system configuration

Task	Storage management utility
Create disk arrays	HP Array Configuration Utility
Create logical disks from the array space	HP Array Configuration Utility
Verify newly created logical disks	Windows Disk Management
Create a volume on the new logical disk	Windows Disk Management

- Create disk arrays—On storage systems with configurable storage, physical disks can be arranged as RAID arrays for fault tolerance and enhanced performance, and then segmented into logical disks of appropriate sizes for particular storage needs. These logical disks then become the volumes that appear as drives on the storage system.

△ CAUTION:

The first two logical drives are configured for the storage system operating system and should not be altered in any manner. If the first two logical drives are altered, the system recovery process may not function properly when using the System Recovery DVD. Do not tamper with the "DON'T ERASE" or local C: volume. These are reserved volumes and must be maintained as they exist.

The fault tolerance level depends on the amount of disks selected when the array was created. A minimum of two disks is required for RAID 0+1 configuration, three disks for a RAID 5 configuration, and four disks for a RAID 6 (ADG) configuration.

- Create logical disks from the array space—Select the desired fault tolerance, stripe size, and size of the logical disk.
- Verify newly created logical disks—Verify that disks matching the newly created sizes are displayed.
- Create a volume on the new logical disk—Select a drive letter and enter a volume label, volume size, allocation unit size, and mount point (if desired).

Storage management utilities

The storage management utilities preinstalled on the storage system include the HP Array Configuration Utility (ACU).

Array management utilities

Storage devices for RAID arrays and LUNs are created and managed using the array management utilities mentioned previously. For HP Smart Arrays use the ACU.

NOTE:

The ACU is used to configure and manage array-based storage. Software RAID-based storage systems use Microsoft Disk Manager to manage storage. You need administrator or root privileges to run the ACU.

Array Configuration Utility

The HP ACU supports the Smart Array controllers and hard drives installed on the storage system.

To open the ACU from the storage system desktop:

NOTE:

If this is the first time that the ACU is being run, you will be prompted to select the Execution Mode for ACU. Selecting Local Application Mode allows you to run the ACU from a Remote Desktop, remote console, or storage system web access mode. Remote service mode allows you to access the ACU from a remote browser.

1. Select **Start > Programs > HP Management Tools > Array Configuration Utility**.
2. If the Execution Mode for ACU is set to Remote Mode, log on to the HP System Management Homepage. The default user name is **administrator** and the default password is **hpinvent**.

To open the ACU in browser mode:

NOTE:

Confirm that the ACU Execution Mode is set to remote service.

1. Open a browser and enter the server name or IP address of the destination server. For example, `http://servername:2301` or `http://192.0.0.1:2301`.
2. Log on to the HP System Management Homepage. The default user name is **administrator** and the default password is **hpinvent**.
3. Click **Array Configuration Utility** on the left side of the window. The ACU opens and identifies the controllers that are connected to the system.

Some ACU guidelines to consider:

- Do not modify the first two logical drives of the storage system; they are configured for the storage system operating system.
- Spanning more than 14 disks with a RAID 5 volume is not recommended.
- Designate spares for RAID sets to provide greater protection against failures.
- RAID sets cannot span controllers.
- A single array can contain multiple logical drives of varying RAID settings.
- Extending and expanding arrays and logical drives is supported.

The *HP Array Configuration Utility User Guide* is available for download at http://www.hp.com/support/manuals.

Disk Management utility

The Disk Management tool is a system utility for managing hard disks and the volumes, or partitions, that they contain. Disk Management is used to initialize disks, create volumes, format volumes with the FAT, FAT32, or NTFS file systems, and create fault-tolerant disk systems. Most disk-related tasks can be performed in Disk Management without restarting the system or interrupting users. Most configuration changes take effect immediately. A complete online help facility is provided with the Disk Management utility for assistance in using the product.

NOTE:

- When the Disk Management utility is accessed through a Remote Desktop connection, this connection can only be used to manage disks and volumes on the server. Using the Remote Desktop connection for other operations during an open session closes the session.
- When closing Disk Management through a Remote Desktop connection, it may take a few moments for the remote session to log off.

Guidelines for managing disks and volumes

- The first two logical drives are configured for the storage system operating system and should not be altered in any manner. If the first two logical drives are altered, the system recovery process may not function properly when using the System Recovery DVD. Do not tamper with the "DON'T ERASE" or local C: volume. These are reserved volumes and must be maintained as they exist.
- HP does not recommend spanning array controllers with dynamic volumes. The use of software RAID-based dynamic volumes is not recommended. Use the array controller instead; it is more efficient.
- Use meaningful volume labels with the intended drive letter embedded in the volume label, if possible. (For example, volume e: might be named "Disk E:.") Volume labels often serve as the only means of identification.
- Record all volume labels and drive letters in case the system needs to be restored.
- When managing basic disks, only the last partition on the disk can be extended unless the disk is changed to dynamic.
- Basic disks can be converted to dynamic, but cannot be converted back to basic without deleting all data on the disk.
- Basic disks can contain up to four primary partitions (or three primary partitions and one extended partition).
- Format drives with a 16 K allocation size for best support of shadow copies, performance, and defragmentation.
- NTFS formatted drives are recommended because they provide the greatest level of support for shadow copies, encryption, and compression.
- Only basic disks can be formatted as FAT or FAT32.
- Read the online Disk Management help found in the utility.

Disk quotas

Disk quotas track and control disk space use in volumes.

NOTE:

To limit the size of a folder or share, see "Quota management" on page 95 .

Configure the volumes on the server to perform the following tasks:

- Prevent further disk space use and log an event when a user exceeds a specified disk space limit.
- Log an event when a user exceeds a specified disk space warning level.

When enabling disk quotas, it is possible to set both the disk quota limit and the disk quota warning level. The disk quota limit specifies the amount of disk space a user is allowed to use. The warning level specifies the point at which a user is nearing his or her quota limit. For example, a user's disk quota limit can be set to 50 megabytes (MB), and the disk quota warning level to 45 MB. In this case, the user can store no more than 50 MB on the volume. If the user stores more than 45 MB on the volume, the disk quota system logs a system event.

In addition, it is possible to specify that users can exceed their quota limit. Enabling quotas and not limiting disk space use is useful to still allow users access to a volume, but track disk space use on a per-user basis. It is also possible to specify whether or not to log an event when users exceed either their quota warning level or their quota limit.

When enabling disk quotas for a volume, volume usage is automatically tracked from that point forward, but existing volume users have no disk quotas applied to them. Apply disk quotas to existing volume users by adding new quota entries on the Quota Entries page.

NOTE:

When enabling disk quotas on a volume, any users with write access to the volume who have not exceeded their quota limit can store data on the volume. The first time a user writes data to a quota-enabled volume, default values for disk space limit and warning level are automatically assigned by the quota system.

For more information about disk quotas, read the online help.

Adding storage

Expansion is the process of adding physical disks to an array that has already been configured. Extension is the process of adding new storage space to an existing logical drive on the same array, usually after the array has been expanded.

Storage growth may occur in three forms:

- Extend unallocated space from the original logical disks or LUNs.
- Alter LUNs to contain additional storage.
- Add new LUNs to the system.

The additional space is then extended through a variety of means, depending on which type of disk structure is in use.

NOTE:

This section addresses only single storage system node configurations. If your server has Windows Storage Server 2008 R2 Enterprise Edition, see the Cluster Administration chapter for expanding and extending storage in a cluster environment.

Expanding storage

Expansion is the process of adding physical disks to an array that has already been configured. The logical drives (or volumes) that exist in the array before the expansion takes place are unchanged, because only the amount of free space in the array changes. The expansion process is entirely independent of the operating system.

NOTE:

See your storage array hardware user documentation for further details about expanding storage on the array.

Extending storage using Windows Storage Utilities

Volume extension grows the storage space of a logical drive. During this process, the administrator adds new storage space to an existing logical drive on the same array, usually after the array has been expanded. An administrator may have gained this new storage space by either expansion or by deleting another logical drive on the same array. Unlike drive expansion, the operating system must be aware of changes to the logical drive size.

You extend a volume to:

- Increase raw data storage
- Improve performance by increasing the number of spindles in a logical drive volume
- Change fault-tolerance (RAID) configurations

For more information about RAID levels, see the *Smart Array Controller User Guide*, or the document titled *Assessing RAID ADG vs. RAID 5 vs. RAID 1+0*. Both are available at the Smart Array controller web page or at http://h18000.www1.hp.com/products/servers/proliantstorage/arraycontrollers/documentation.html.

Extend volumes using Disk Management

The Disk Management snap-in provides management of hard disks, volumes or partitions. It can be used to extend a dynamic volume only.

NOTE:

Disk Management cannot be used to extend basic disk partitions.

Guidelines for extending a dynamic volume:

- Use the Disk Management utility.
- You can extend a volume only if it does not have a file system or if it is formatted NTFS.
- You cannot extend volumes formatted using FAT or FAT32.

- You cannot extend striped volumes, mirrored volumes, or RAID 5 volumes.

For more information, see the Disk Management online help.

Volume shadow copies

NOTE:

Select storage systems can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses using shadow copies in a non-clustered environment.

The Volume Shadow Copy Service provides an infrastructure for creating point-in-time snapshots (shadow copies) of volumes. Shadow Copy supports 64 shadow copies per volume.

A shadow copy contains previous versions of the files or folders contained on a volume at a specific point in time. While the shadow copy mechanism is managed at the server, previous versions of files and folders are only available over the network from clients, and are seen on a per folder or file level, and not as an entire volume.

The shadow copy feature uses data blocks. As changes are made to the file system, the Shadow Copy Service copies the original blocks to a special cache file to maintain a consistent view of the file at a particular point in time. Because the snapshot only contains a subset of the original blocks, the cache file is typically smaller than the original volume. In the snapshot's original form, it takes up no space because blocks are not moved until an update to the disk occurs.

By using shadow copies, a storage system can maintain a set of previous versions of all files on the selected volumes. End users access the file or folder by using a separate client add-on program, which enables them to view the file in Windows Explorer. Accessing previous versions of files, or shadow copies, enables users to:

- Recover files that were accidentally deleted. Previous versions can be opened and copied to a safe location.
- Recover from accidentally overwriting a file. A previous version of that file can be accessed.
- Compare several versions of a file while working. Use previous versions to compare changes between two versions of a file.

Shadow copies cannot replace the current backup, archive, or business recovery system, but they can help to simplify restore procedures. Because a snapshot only contains a portion of the original data blocks, shadow copies cannot protect against data loss due to media failures. However, the strength of snapshots is the ability to instantly recover data from shadow copies, reducing the number of times needed to restore data from tape.

Shadow copy planning

Before setup is initiated on the server and the client interface is made available to end users, consider the following:

- From what volume will shadow copies be taken?
- How much disk space should be allocated for shadow copies?
- Will separate disks be used to store shadow copies?
- How frequently will shadow copies be made?

Identifying the volume

Shadow copies are taken for a complete volume, but not for a specific directory. Shadow copies work best when the server stores user files, such as documents, spreadsheets, presentations, graphics, or database files.

NOTE:

Shadow copies should not be used to provide access to previous versions of application or e-mail databases.

Shadow copies are designed for volumes that store user data such as home directories and My Documents folders that are redirected by using Group Policy or other shared folders in which users store data.

Shadow copies work with compressed or encrypted files and retain whatever permissions were set on the files when the shadow copies were taken. For example, if a user is denied permission to read a file, that user would not be able to restore a previous version of the file, or be able to read the file after it has been restored.

Although shadow copies are taken for an entire volume, users must use shared folders to access shadow copies. Administrators on the local server must also specify the `\\servername\sharename` path to access shadow copies. If administrators or end users want to access a previous version of a file that does not reside in a shared folder, the administrator must first share the folder.

NOTE:

Shadow copies are available only on NTFS, not FAT or FAT32 volumes.

Files or folders that are recorded by using Shadow Copy appear static, even though the original data is changing.

Allocating disk space

When determining the amount of space to allocate for storing shadow copies, consider both the number and size of files that are being copied, as well as the frequency of changes between copies. For example, 100 files that only change monthly require less storage space than 10 files that change daily. If the frequency of changes to each file is greater than the amount of space allocated to storing shadow copies, no shadow copy is created.

Administrators should also consider user expectations of how many versions they will want to have available. End users might expect only a single shadow copy to be available, or they might expect three days or three weeks worth of shadow copies. The more shadow copies users expect, the more storage space administrators must allocate for storing them.

Setting the limit too low also affects backup programs that use shadow copy technology because these programs are also limited to using the amount of disk space specified by administrators.

NOTE:

Regardless of the volume space that is allocated for shadow copies, there is a maximum of 64 shadow copies for any volume. When the 65th shadow copy is taken, the oldest shadow copy is purged.

The minimum amount of storage space that can be specified is 350 megabytes (MB). The default storage size is 10 percent of the source volume (the volume being copied). If the shadow copies are stored on a separate volume, change the default to reflect the space available on the *storage* volume instead of the *source* volume. Remember that when the storage limit is reached, older versions of the shadow copies are deleted and cannot be restored.

△ CAUTION:

To change the storage volume, shadow copies must be deleted. The existing file change history that is kept on the original storage volume is lost. To avoid this problem, verify that the storage volume that is initially selected is large enough.

Identifying the storage area

To store the shadow copies of another volume on the same file server, a volume can be dedicated on separate disks. For example, if user files are stored on `H:\`, another volume such as `S:\`can be used to store the shadow copies. Using a separate volume on separate disks provides better performance and is recommended for heavily used storage systems.

If a separate volume will be used for the storage area (where shadow copies are stored), the maximum size should be changed to **No Limit** to reflect the space available on the storage area volume instead of the source volume (where the user files are stored).

Disk space for shadow copies can be allocated on either the same volume as the source files or a different volume. There is a trade-off between ease of use and maintenance versus performance and reliability that the system administrator must consider.

By keeping the shadow copy on the same volume, there is a potential gain in ease of setup and maintenance; however, there may be a reduction in performance and reliability.

△ CAUTION:

If shadow copies are stored on the same volume as the user files, note that a burst of disk input/output (I/O) can cause all shadow copies to be deleted. If the sudden deletion of shadow copies is unacceptable to administrators or end users, it is best to use a separate volume on separate disks to store shadow copies.

Determining creation frequency

The more frequently shadow copies are created, the more likely that end users will get the version that they want. However, with a maximum of 64 shadow copies per volume, there is a trade-off between the frequency of making shadow copies and the amount of time that the earlier files will be available.

By default, the storage system creates shadow copies at 0700 and 1200, Monday through Friday. However, these settings are easily modified by the administrator so that the shadow copy schedule can better accommodate end user needs.

Shadow copies and drive defragmentation

When running Disk Defragmenter on a volume with shadow copies activated, all or some of the shadow copies may be lost, starting with the oldest shadow copies.

If defragmenting volumes on which shadow copies are enabled, use a cluster (or allocation unit) size of 16 KB or larger. Using this allocation unit size reduces the number of copy outs occurring on the snapshot. Otherwise, the number of changes caused by the defragmentation process can cause shadow copies to be deleted faster than expected. Note, however, that NTFS compression is supported only if the cluster size is 4 KB or smaller.

NOTE:

To check the cluster size of a volume, use the `fsutil fsinfo ntfsinfo` command. To change the cluster size on a volume that contains data, back up the data on the volume, reformat it using the new cluster size, and then restore the data.

Mounted drives

A mounted drive is a local volume attached to an empty folder (called a mount point) on an NTFS volume. When enabling shadow copies on a volume that contains mounted drives, the mounted drives are not included when shadow copies are taken. In addition, if a mounted drive is shared and shadow copies are enabled on it, users cannot access the shadow copies if they traverse from the host volume (where the mount point is stored) to the mounted drive.

For example, assume there is a folder `F:\data\users`, and the `Users` folder is a mount point for `G:\`. If shadow copies are enabled on both `F:\` and `G:\`, `F:\data` is shared as `\\server1\data`, and `G:\data\users` is shared as `\\server1\users`. In this example, users can access previous versions of `\\server1\data` and `\\server1\users` but not `\\server1\data\users`.

Managing shadow copies

The vssadmin tool provides a command line capability to create, list, resize, and delete volume shadow copies.

The system administrator can make shadow copies available to end users through a feature called "Shadow Copies for Shared Folders." The administrator uses the Properties menu (see Figure 29) to turn on the Shadow Copies feature, select the volumes to be copied, and determine the frequency with which shadow copies are made.

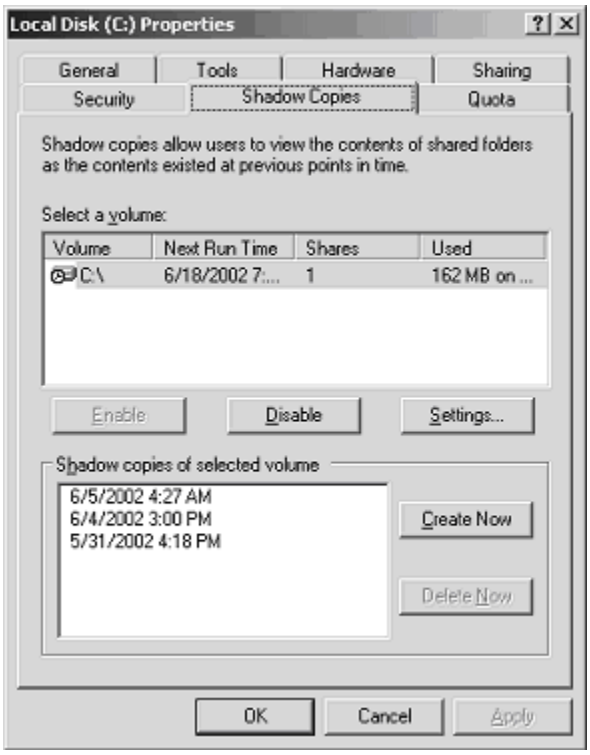

Figure 29 System administrator view of Shadow Copies for Shared Folders

The shadow copy cache file

The default shadow copy settings allocate 10 percent of the source volume being copied (with a minimum of 350 MB), and store the shadow copies on the same volume as the original volume. (See Figure 30). The cache file is located in a hidden protected directory titled "System Volume Information" off of the root of each volume for which shadow copy is enabled.

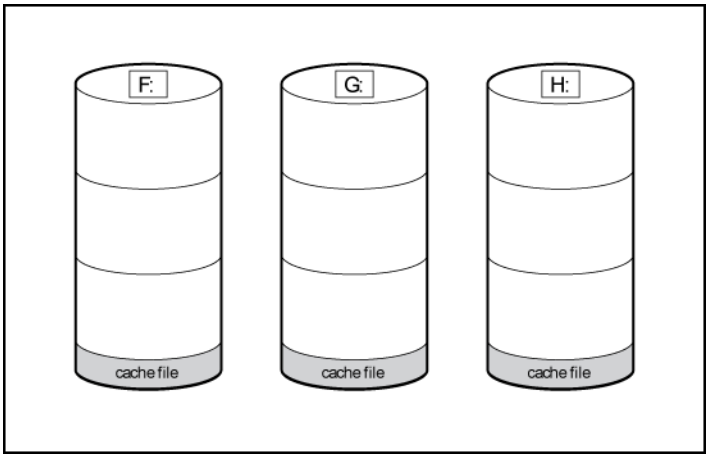

Figure 30 Shadow copies stored on a source volume

The cache file location can be altered to reside on a dedicated volume separate from the volumes containing files shares. (See Figure 31).

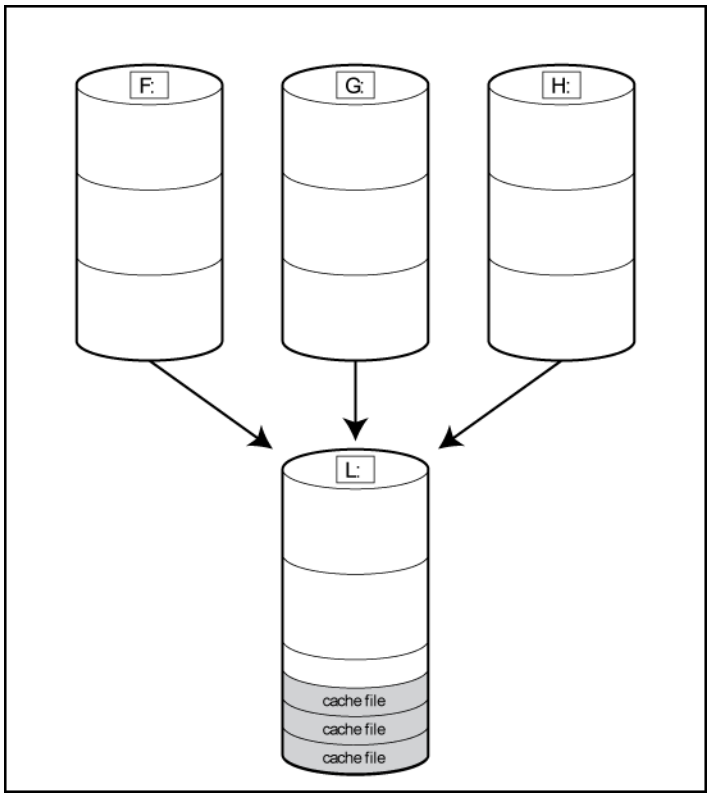

Figure 31 Shadow copies stored on a separate volume

The main advantage to storing shadow copies on a separate volume is ease of management and performance. Shadow copies on a source volume must be continually monitored and can consume space designated for file sharing. Setting the limit too high takes up valuable storage space. Setting the limit too low can cause shadow copies to be purged too soon, or not created at all. By storing shadow copies on a separate volume space, limits can generally be set higher, or set to No Limit. See the online help for instructions on altering the cache file location.

△ CAUTION:

If the data on the separate volume L: is lost, the shadow copies cannot be recovered.

Enabling and creating shadow copies

Enabling shadow copies on a volume automatically results in several actions:

- Creates a shadow copy of the selected volume.
- Sets the maximum storage space for the shadow copies.
- Schedules shadow copies to be made at 7 a.m. and 12 noon on weekdays.

NOTE:

Creating a shadow copy only makes one copy of the volume; it does not create a schedule.

NOTE:

After the first shadow copy is created, it cannot be relocated. Relocate the cache file by altering the cache file location under Properties prior to enabling shadow copy. See "Viewing shadow copy properties" on page 79.

Viewing a list of shadow copies

To view a list of shadow copies on a volume:

1. Access Disk Management.
2. Select the volume or logical drive, then right-click on it.
3. Select **Properties**.
4. Select **Shadow Copies** tab.

All shadow copies are listed, sorted by the date and time they were created.

NOTE:

It is also possible to create new shadow copies or delete shadow copies from this page.

Set schedules

Shadow copy schedules control how frequently shadow copies of a volume are made. There are a number of factors that can help determine the most effective shadow copy schedule for an organization. These include the work habits and locations of the users. For example, if users do not all live in the same time zone, or they work on different schedules, it is possible to adjust the daily shadow copy schedule to allow for these differences.

Do not schedule shadow copies more frequently than once per hour.

NOTE:

When deleting a shadow copy schedule, that action has no effect on existing shadow copies.

Viewing shadow copy properties

The Shadow Copy Properties page lists the number of copies, the date and time the most recent shadow copy was made, and the maximum size setting.

NOTE:

For volumes where shadow copies do not exist currently, it is possible to change the location of the cache file. Managing the cache files on a separate disk is recommended.

CAUTION:

Use caution when reducing the size limit for all shadow copies. When the size is set to less than the total size currently used for all shadow copies, enough shadow copies are deleted to reduce the total size to the new limit. A shadow copy cannot be recovered after it has been deleted.

Redirecting shadow copies to an alternate volume

IMPORTANT:

Shadow copies must be initially disabled on the volume before redirecting to an alternate volume. If shadow copies are enabled and you disable them, a message appears informing you that all existing shadow copies on the volume will be permanently deleted.

To redirect shadow copies to an alternate volume:

1. Access Disk Management.
2. Select the volume or logical drive, then right-click on it.
3. Select **Properties**.
4. Select the **Shadow Copies** tab.
5. Select the volume that you want to redirect shadow copies from and ensure that shadow copies are disabled on that volume; if enabled, click **Disable**.
6. Click **Settings**.
7. In the **Located on this volume** field, select an available alternate volume from the list.

NOTE:

To change the default shadow copy schedule settings, click **Schedule**.

8. Click **OK**.
9. On the **Shadow Copies** tab, ensure that the volume is selected, and then click **Enable**.

Shadow copies are now scheduled to be made on the alternate volume.

Disabling shadow copies

When shadow copies are disabled on a volume, all existing shadow copies on the volume are deleted as well as the schedule for making new shadow copies.

CAUTION:

When the Shadow Copies Service is disabled, all shadow copies on the selected volumes are deleted. Once deleted, shadow copies cannot be restored.

Managing shadow copies from the storage system desktop

To access shadow copies from the storage system desktop:

The storage system desktop can be accessed by using Remote Desktop to manage shadow copies.

1. On the storage system desktop, double-click **My Computer**.
2. Right-click the volume name, and select **Properties**.
3. Click the **Shadow Copies** tab. See Figure 32.

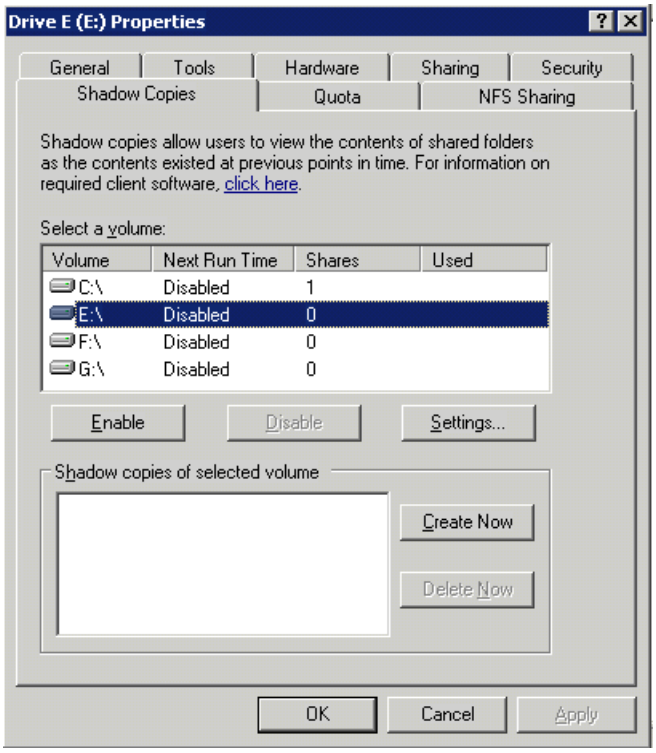

Figure 32 Accessing shadow copies from My Computer

Shadow Copies for Shared Folders

Shadow copies are accessed over the network by supported clients and protocols. There are two sets of supported protocols, SMB and NFS. All other protocols are not supported; this includes HTTP, FTP, AppleTalk, and NetWare Shares. For SMB support, a client-side application denoted as Shadow Copies for Shared Folders is required. The client-side application is currently only available for Windows XP and Windows 2000 SP3+.

No additional software is required to enable UNIX users to independently retrieve previous versions of files stored on NFS shares.

NOTE:

Shadow Copies for Shared Folders supports retrieval only of shadow copies of network shares. It does not support retrieval of shadow copies of local folders.

NOTE:

Shadow Copies for Shared Folders clients are not available for HTTP, FTP, AppleTalk, or NetWare shares. Consequently, users of these protocols cannot use Shadow Copies for Shared Folders to independently retrieve previous versions of their files. However, administrators can take advantage of Shadow Copies for Shared Folders to restore files for these users.

SMB shadow copies

Windows users can independently access previous versions of files stored on SMB shares by using the Shadow Copies for Shared Folders client. After the Shadow Copies for Shared Folders client is installed on the user's computer, the user can access shadow copies for a share by right-clicking on the share to open its Properties window, clicking the **Previous Versions** tab, and then selecting the desired shadow copy. Users can view, copy, and restore all available shadow copies.

Shadow Copies for Shared Folders preserves the permissions set in the access control list (ACL) of the original folders and files. Consequently, users can only access shadow copies for shares to which they have access. In other words, if a user does not have access to a share, he also does not have access to the share's shadow copies.

The Shadow Copies for Shared Folders client pack installs a **Previous Versions** tab in the **Properties** window of files and folders on network shares.

Users access shadow copies with Windows Explorer by selecting **View**, **Copy**, or **Restore** from the **Previous Versions** tab. (See Figure 33). Both individual files and folders can be restored.

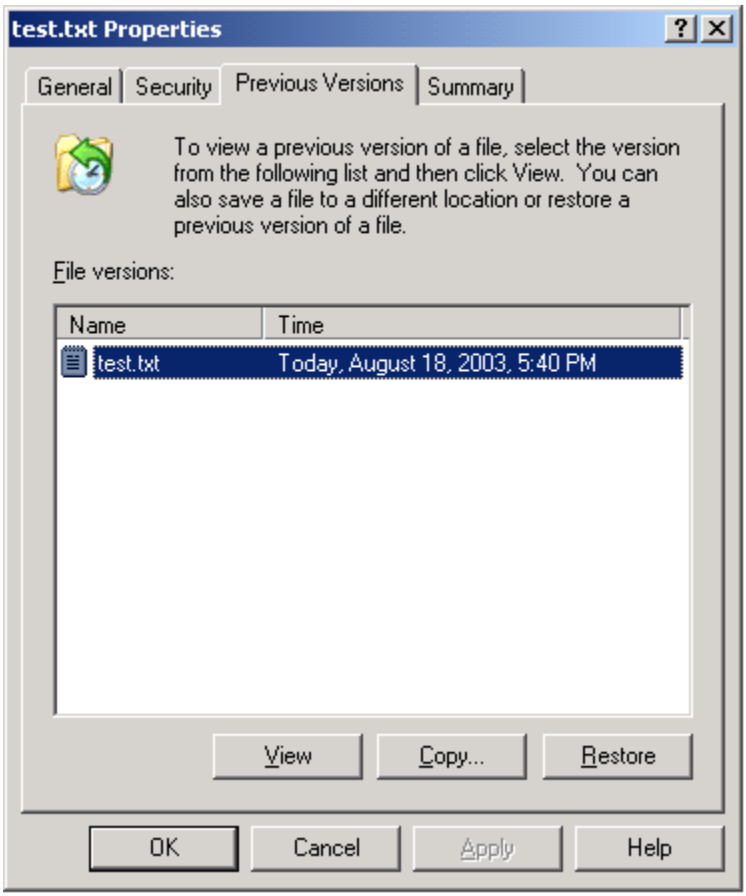

Figure 33 Client GUI

When users view a network folder hosted on the storage system for which shadow copies are enabled, old versions (prior to the snapshot) of a file or directory are available. Viewing the properties of the file or folder presents users with the folder or file history—a list of read-only, point-in-time copies of the file or folder contents that users can then open and explore like any other file or folder. Users can view files in the folder history, copy files from the folder history, and so on.

NFS shadow copies

UNIX users can independently access previous versions of files stored on NFS shares via the NFS client; no additional software is required. Server for NFS exposes each of a share's available shadow copies as a pseudo-subdirectory of the share. Each of these pseudo-subdirectories is displayed in exactly the same way as a regular subdirectory is displayed.

The name of each pseudo-subdirectory reflects the creation time of the shadow copy, using the format .@GMT-YYYY.MM.DD-HH:MM:SS. To prevent common tools from needlessly enumerating the pseudo-subdirectories, the name of each pseudo-subdirectory begins with the dot character, thus rendering it hidden.

The following example shows an NFS share named “NFSShare” with three shadow copies, taken on April 27, 28, and 29 of 2003 at 4 a.m.

NFSShare

.@GMT-2003.04.27-04:00:00

.@GMT-2003.04.28-04:00:00

.@GMT-2003.04.29-04:00:00

Access to NFS shadow copy pseudo-subdirectories is governed by normal access-control mechanisms using the permissions stored in the file system. Users can access only those shadow copies to which they have read access at the time the shadow copy is taken. To prevent users from modifying shadow copies, all pseudo-subdirectories are marked read-only, regardless of the user's ownership or access rights, or the permissions set on the original files.

Server for NFS periodically polls the system for the arrival or removal of shadow copies and updates the root directory view accordingly. Clients then capture the updated view the next time they issue a directory read on the root of the share.

Recovery of files or folders

There are three common situations that may require recovery of files or folders:

- Accidental file deletion, the most common situation
- Accidental file replacement, which may occur if a user selects Save instead of Save As
- File corruption

It is possible to recover from all of these scenarios by accessing shadow copies. There are separate steps for accessing a file compared to accessing a folder.

Recovering a deleted file or folder

To recover a deleted file or folder within a folder:

1. Access to the folder where the deleted file was stored.
2. Position the cursor over a blank space in the folder. If the cursor hovers over a file, that file is selected.
3. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
4. Select the version of the folder that contains the file before it was deleted, and then click **View**.
5. View the folder and select the file or folder to recover. The view may be navigated multiple folders deep.
6. Click **Restore** to restore the file or folder to its original location. Click **Copy...** to allow the placement of the file or folder to a new location.

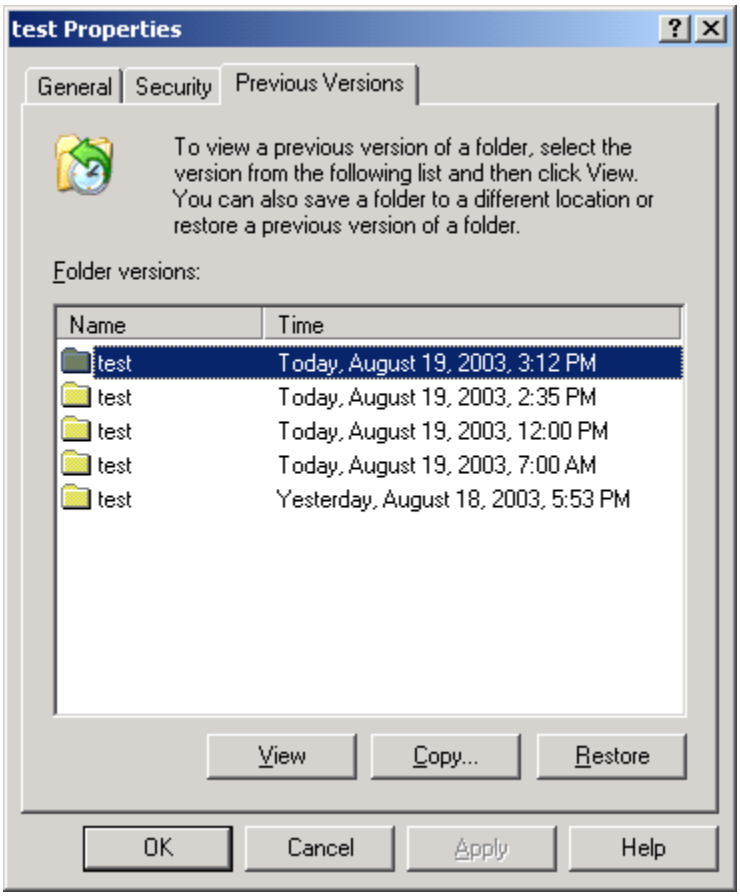

Figure 34 Recovering a deleted file or folder

Recovering an overwritten or corrupted file

Recovering an overwritten or corrupted file is easier than recovering a deleted file because the file itself can be right-clicked instead of the folder. To recover an overwritten or corrupted file:

1. Right-click the overwritten or corrupted file, and then click **Properties**.
2. Click **Previous Versions**.
3. To view the old version, click **View**. To copy the old version to another location, click **Copy...** to replace the current version with the older version, click **Restore**.

Recovering a folder

To recover a folder:

1. Position the cursor so that it is over a blank space in the folder to be recovered. If the cursor hovers over a file, that file is selected.
2. Right-click, select **Properties** from the bottom of the menu, and then click the **Previous Versions** tab.
3. Click either **Copy...** or **Restore**.

   Clicking **Restore** enables the user to recover everything in that folder as well as all subfolders. Clicking **Restore** does not delete any files.

Backup and shadow copies

Shadow copies are only available on the network via the client application, and only at a file or folder level as opposed to the entire volume. Hence, the standard backup associated with a volume backup will not work to back up the previous versions of the file system. To answer this particular issue, shadow copies are available for backup in two situations. If the backup software in question supports the use of shadow copies and can communicate with underlying block device, it is supported, and the previous version of the file system will be listed in the backup application as a complete file system snapshot. If the built-in backup application NTbackup is used, the backup software forces a snapshot, and then uses the snapshot as the means for backup. The user is unaware of this activity and it is not self-evident although it does address the issue of open files.

Shadow Copy Transport

Shadow Copy Transport provides the ability to transport data on a Storage Area Network (SAN). With a storage array and a VSS-aware hardware provider, it is possible to create a shadow copy on one server and import it on another server. This process, essentially "virtual" transport, is accomplished in a matter of minutes, regardless of the size of the data.

A shadow copy transport can be used for a number of purposes, including:

- Tape backups

  An alternative to traditional backup to tape processes is transport of shadow copies from the production server onto a backup server, where they can then be backed up to tape. Like the other two alternatives, this option removes backup traffic from the production server. While some backup applications might be designed with the hardware provider software that enables transport, others are not. The administrator should determine whether or not this functionality is included in the backup application.

- Data mining

  The data in use by a particular production server is often useful to different groups or departments within an organization. Rather than add additional traffic to the production server, a shadow copy of the data can be made available through transport to another server. The shadow copy can then be processed for different purposes, without any performance impact on the original server.

The transport process is accomplished through a series of DISKRAID command steps:

1. Create a shadow copy of the source data on the source server (read-only).
2. Mask off (hide) the shadow copy from the source server.
3. Unmask the shadow copy to a target server.
4. Optionally, clear the read-only flags on the shadow copy.

The data is now ready to use.

Folder and share management

The storage system supports several file-sharing protocols, including DFS, NFS, FTP, HTTP, and Microsoft SMB. This section discusses overview information as well as procedures for the setup and management of the file shares for the supported protocols. Security at the file level and at the share level is also discussed.

NOTE:

Select servers can be deployed in a clustered or non-clustered configuration. This section discusses share setup for a non-clustered deployment.

Folder management

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders eases the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Folders can be managed using Server Manager. Tasks include:

- Accessing a specific volume or folder
- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share for a volume or folder
- Managing shares for a volume or folder

Managing file-level permissions

Security at the file level is managed using Windows Explorer.

File level security includes settings for permissions, ownership, and auditing for individual files.

To enter file permissions:

1. Using Windows Explorer, access the folder or file that needs to be changed, and then right-click the folder.

2. Click **Properties**, and then click the **Security** tab.

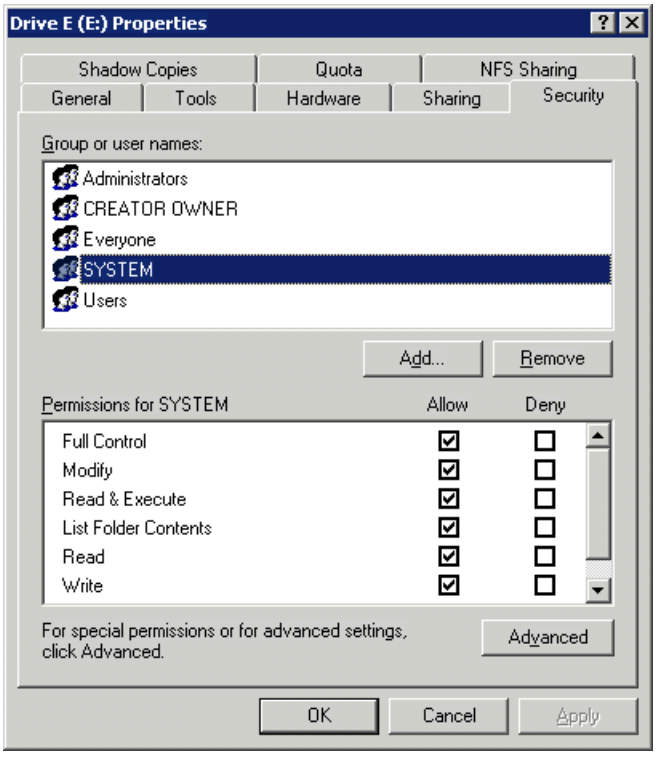

Figure 35 Properties dialog box, Security tab

Several options are available on the **Security** tab:

- To add users and groups to the permissions list, click **Add**. Follow the dialog box instructions.
- To remove users and groups from the permissions list, highlight the desired user or group, and then click **Remove**.
- The center section of the **Security** tab lists permission levels. When new users or groups are added to the permissions list, select the appropriate boxes to configure the common file-access levels.

3. To modify ownership of files, or to modify individual file access level permissions, click **Advanced**.

   Figure 36 illustrates the properties available on the **Advanced Security Settings** dialog box.

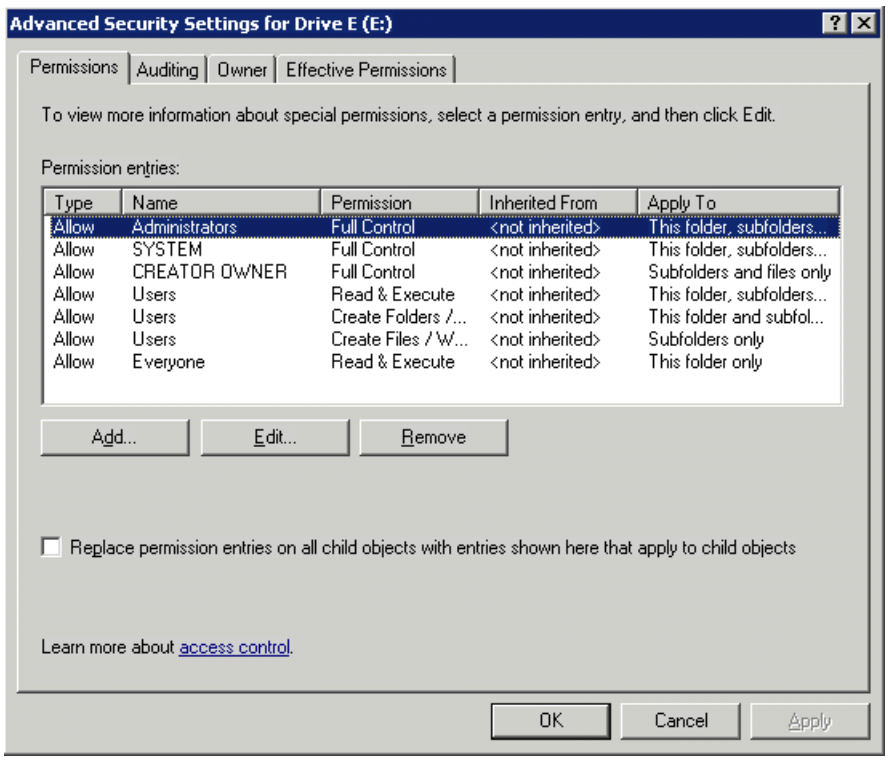

Figure 36 Advanced Security settings dialog box, Permissions tab

Other functionality available in the **Advanced Security Settings** dialog box is illustrated in Figure 36 and includes:

- Add a new user or group—Click **Add**, and then follow the dialog box instructions.
- Remove a user or group— Click **Remove**.
- Replace permission entries on all child objects with entries shown here that apply to child objects—This allows all child folders and files to inherit the current folder permissions by default.
- Modify specific permissions assigned to a particular user or group—Select the desired user or group, and then click **Edit**.

4. Enable or disable permissions by selecting the **Allow** box to enable permission or the **Deny** box to disable permission. If neither box is selected, permission is automatically disabled. Figure 37 illustrates the **Edit** screen and some of the permissions.

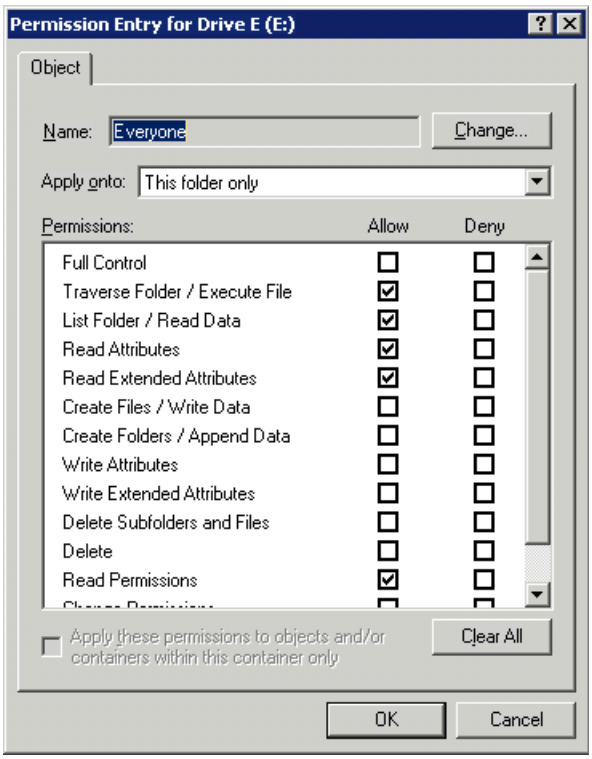

Figure 37 User or group Permission Entry dialog box

Another area of the **Advanced Security Settings** is the **Auditing** tab. Auditing allows you to set rules for the auditing of access, or attempted access, to files or folders. Users or groups can be added, deleted, viewed, or modified through the **Advanced Security Settings Auditing** tab.

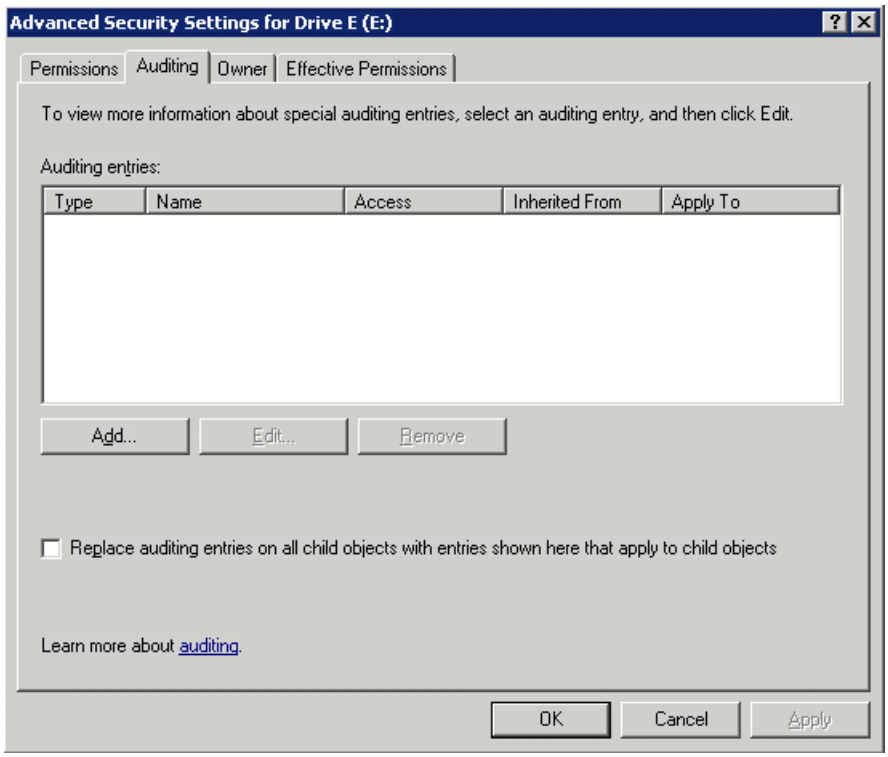

Figure 38 Advanced Security Settings dialog box, Auditing tab

5. Click **Add** to display the Select User or Group dialog box.

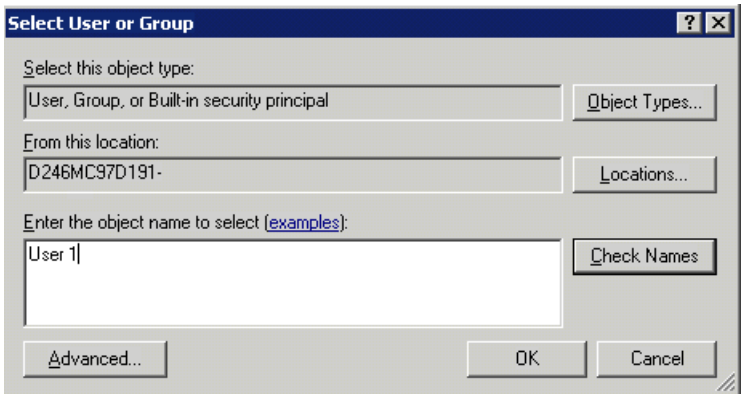

Figure 39 Select User or Group dialog box

NOTE:

Click Advanced to search for users or groups.

6. Select the user or group.

7. Click **OK**.

   The **Auditing Entry** dialog box is displayed.

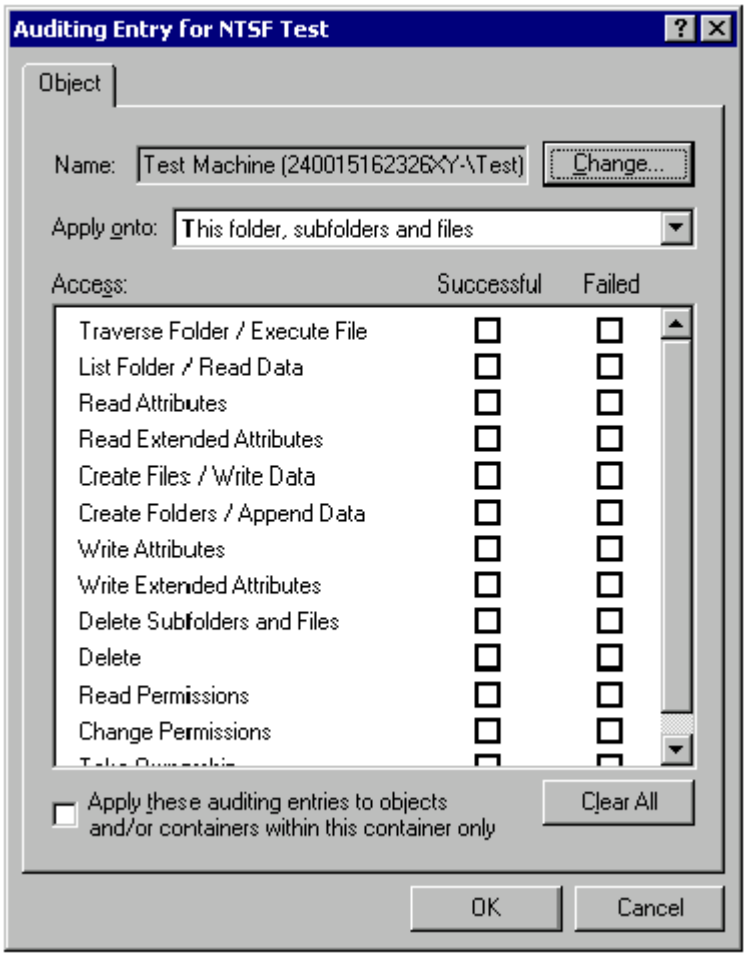

Figure 40 Auditing Entry dialog box for folder name NTFS Test

8. Select the desired **Successful** and **Failed** audits for the user or group.
9. Click **OK**.

NOTE:

Auditing must be enabled to configure this information. Use the local Computer Policy Editor to configure the audit policy on the storage system.

The **Owner** tab allows taking ownership of files. Typically, administrators use this area to take ownership of files when the file ACL is incomplete or corrupt. By taking ownership, you gain access to the files, and then manually apply the appropriate security configurations.

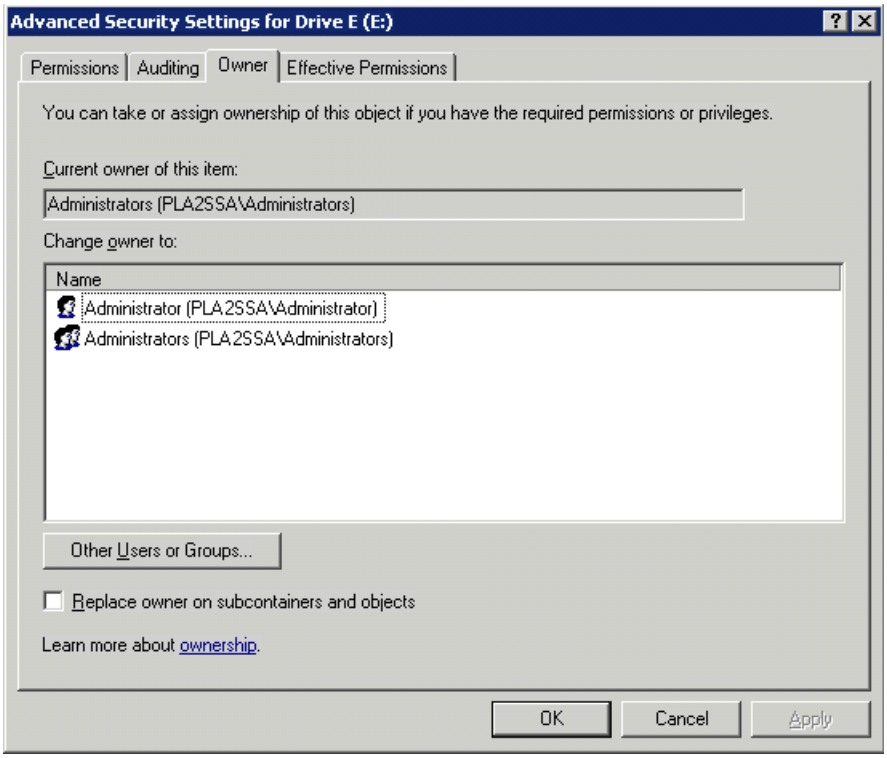

Figure 41 Advanced Security Settings dialog box, Owner tab

The current owner of the file or folder is listed at the top of the screen. To take ownership:

1. Click the appropriate user or group in the **Change owner to** list.
2. If it is also necessary to take ownership of subfolders and files, enable the **Replace owner on subcontainers and objects** box.
3. Click **OK**.

Share management

There are several ways to set up and manage shares. Methods include using Windows Explorer, a command line interface, or Server Manger.

NOTE:

Select servers can be deployed in a clustered as well as a non-clustered configuration. This chapter discusses share setup for a non-clustered deployment.

As previously mentioned, the file-sharing security model of the storage system is based on the NTFS file-level security model. Share security seamlessly integrates with file security. In addition to discussing share management, this section discusses share security.

Share considerations

Planning the content, size, and distribution of shares on the storage system can improve performance, manageability, and ease of use.

The content of shares should be carefully chosen to avoid two common pitfalls: either having too many shares of a very specific nature, or of having very few shares of a generic nature. For example, shares for general use are easier to set up in the beginning, but can cause problems later. Frequently, a better approach is to create separate shares with a specific purpose or group of users in mind. However, creating too many shares also has its drawbacks. For example, if it is sufficient to create a single share for user home directories, create a “homes” share rather than creating separate shares for each user.

By keeping the number of shares and other resources low, the performance of the storage system is optimized. For example, instead of sharing out each individual user's home directory as its own share, share out the top-level directory and let the users map personal drives to their own subdirectory.

Defining Access Control Lists

The Access Control List (ACL) contains the information that dictates which users and groups have access to a share, as well as the type of access that is permitted. Each share on an NTFS file system has one ACL with multiple associated user permissions. For example, an ACL can define that User1 has read and write access to a share, User2 has read only access, and User3 has no access to the share. The ACL also includes group access information that applies to every user in a configured group. ACLs are also referred to as permissions.

Integrating local file system security into Windows domain environments

ACLs include properties specific to users and groups from a particular workgroup server or domain environment. In a multidomain environment, user and group permissions from several domains can apply to files stored on the same device. Users and groups local to the storage system can be given access permissions to shares managed by the device. The domain name of the storage system supplies the context in which the user or group is understood. Permission configuration depends on the network and domain infrastructure where the server resides.

File-sharing protocols (except NFS) supply a user and group context for all connections over the network. (NFS supplies a machine-based context.) When new files are created by those users or machines, the appropriate ACLs are applied.

Configuration tools provide the ability to share permissions out to clients. These shared permissions are propagated into a file system ACL, and when new files are created over the network, the user creating the file becomes the file owner. In cases where a specific subdirectory of a share has different permissions from the share itself, the NTFS permissions on the subdirectory apply instead. This method results in a hierarchical security model where the network protocol permissions and the file permissions work together to provide appropriate security for shares on the device.

NOTE:

Share permissions and file-level permissions are implemented separately. It is possible for files on a file system to have different permissions from those applied to a share. When this situation occurs, the file-level permissions override the share permissions.

Comparing administrative (hidden) and standard shares

CIFS supports both administrative shares and standard shares.

- Administrative shares are shares with a last character of $. Administrative shares are not included in the list of shares when a client browses for available shares on a CIFS server.

- Standard shares are shares that do not end in a $ character. Standard shares are listed whenever a CIFS client browses for available shares on a CIFS server.

The storage system supports both administrative and standard CIFS shares. To create an administrative share, end the share name with the $ character when setting up the share. Do not type a $ character at the end of the share name when creating a standard share.

Managing shares

Shares can be managed using Server Manager. Tasks include:

- Creating a new share
- Deleting a share
- Modifying share properties
- Publishing in DFS

NOTE:

These functions can operate in a cluster on select servers, but should only be used for non-cluster-aware shares. Use Failover Cluster Management to manage shares for a cluster. The page will display cluster share resources.

CAUTION:

Before deleting a share, warn all users to exit that share and confirm that no one is using that share.

File Server Resource Manager

File Server Resource Manager (FSRM) is a suite of tools that allows administrators to understand, control, and manage the quantity and type of data stored on their servers. Some of the tasks you can perform are:

- Quota management
- File screening management
- Storage reports

Server Manager provides access to FSRM tasks.

For procedures and methods beyond what are described below, see the online help.

Quota management

On the Quota Management node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Create quotas to limit the space allowed for a volume or folder and generate notifications when the quota limits are approached or exceeded.
- Generate auto quotas that apply to all existing folders in a volume or folder, as well as to any new subfolders created in the future.
- Define quota templates that can be easily applied to new volumes or folders and that can be used across an organization.

File screening management

On the File Screening Management node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Create file screens to control the types of files that users can save and to send notifications when users attempt to save blocked files.
- Define file screening templates that can be easily applied to new volumes or folders and that can be used across an organization.
- Create file screening exceptions that extend the flexibility of the file screening rules.

Storage reports

On the Storage Reports node of the File Server Resource Manager snap-in, you can perform the following tasks:

- Schedule periodic storage reports that allow you to identify trends in disk usage.
- Monitor attempts to save unauthorized files for all users or a selected group of users.
- Generate storage reports instantly.

Other Windows disk and data management tools

When you install certain tools, such as Windows Support Tools or Windows Resource Kit Tools, information about these tools might appear in Help and Support Center. To see the tools that are available to you, look in the Help and Support Center under **Support Tasks**, click **Tools**, and then click **Tools by Category**.

NOTE:

The Windows Support Tools and Windows Resource Kit Tools, including documentation for these tools, are available in English only. If you install them on a non-English language operating system or on an operating system with a Multilingual User Interface Pack (MUI), you see English content mixed with non-English content in Help and Support Center. To see the tools that are available to you, click **Start**, click **Help and Support Center**, and then, under **Support Tasks**, click **Tools**.

Additional information and references for file services

Backup

HP recommends that you back up the print server configuration whenever a new printer is added to the network and the print server configuration is modified.

HP StorageWorks Library and Tape Tools

HP StorageWorks Library and Tape Tools (L&TT) provides functionality for firmware downloads, verification of device operation, maintenance procedures, failure analysis, corrective service actions, and some utility functions. It also provides seamless integration with HP hardware support by generating and e-mailing support tickets that deliver a snapshot of the storage system.

For more information, and to download the utility, see the StorageWorks L&TT web site at http://h18006.www1.hp.com/products/storageworks/ltt.

Antivirus

The server should be secured by installing the appropriate antivirus software.anything

7 Troubleshooting, servicing, and maintenance

Troubleshooting the storage system

The "Support and troubleshooting" task at the HP Support & Drivers web site (http://www.hp.com/go/support) can be used to troubleshoot problems with the storage system. After entering the storage system name and designation (for example, ML110 G5 storage system) or component information (for example, Array Configuration Utility), use the following links for troubleshooting information:

- Download drivers and software—This area provides drivers and software for your operating system.
- Troubleshoot a problem—This area provides a listing of customer notices, advisories, and bulletins applicable for the product or component.
- Manuals—This area provides the latest user documentation applicable to the product or component. User guides can be a useful source for troubleshooting information. For most storage system hardware platforms, the following ProLiant server manuals may be useful for troubleshooting assistance:
  - **HP ProLiant Server User Guide or HP ProLiant Server Maintenance and Service Guide**.
    These guides contain specific troubleshooting information for the server.
  - **HP ProLiant Servers Troubleshooting Guide**
    The guide provides common procedures and solutions for many levels of troubleshooting with a ProLiant server. The guide is available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00300504/c00300504.pdf.

IMPORTANT:

Some troubleshooting procedures found in ProLiant server guides may not apply to the HP StorageWorks P4000 G2 Unified NAS Gateway. If necessary, check with your HP Support representative for further assistance.

For software related components and issues, online help or user guide documentation may offer troubleshooting assistance. The release notes for the storage system product line is updated frequently. The document contains issues and workarounds to a number of categories for the storage systems.

Known issues and workarounds for the storage system products and the service release are addressed in release notes. To view the latest release notes, go to http://www.hp.com/go/nas, select your product family, product model, click **Support for your product**, and then click **Manuals**.

WEBES (Web Based Enterprise Services)

WEBES is a tool suite aimed at preventing or reducing your system's down time. The tool suite has the following components:

- CCAT (Computer Crash Analysis Tool)
- SEA (System Event Analyzer)

If you have a warranty or service contract with HP you are entitled to these tools free of charge. You must, however, upgrade the tools at least once a year because the software expires after one year. For more information about WEBES, see http://h18023.www1.hp.com/support/svctools/webes/.

To install WEBES on your storage system, run the setup executable located in the `C:\hpnas\Components\WEBES` folder.

Maintenance and service

HP provides specific documentation for maintaining and servicing your storage system and offers a customer self repair program.

Maintenance updates

Regular updates to the storage system are supplied on the HP StorageWorks Service Release DVD. The Service Release DVD can be obtained at http://www.software.hp.com.

Individual updates for each product are available for download from the HP Support web site at http://h18023.www1.hp.com/support/selfrepair/na/replace_part.asp.

System updates

System updates to the hardware (BIOS, firmware, drivers), critical updates, and hotfixes for the operating system and other related software updates are bundled on the Service Release DVD.

Firmware updates

Firmware is software that is stored in Read-Only Memory (ROM). Firmware is responsible for the behavior of the system when it is first switched on and for passing control of the server to the operating system. When referring to the firmware on the system board of the server, it is called the System ROM or the BIOS. When referring to the firmware on another piece of hardware configured in the server, it is called Option ROM. Storage systems have hard drives, Smart Array Controllers, Remote Insight Lights-Out Edition (RILOE), Remote Insight Lights-Out Edition II (RILOE II) and Integrated Lights-Out options that have firmware that can be updated.

It is important to update the firmware (also called "flashing the ROM") as part of regular server maintenance. In addition, checking for specific firmware updates in between regular updates helps to keep the server performing optimally. HP recommends checking for a firmware update before sending a part back to HP for replacement.

Certificate of Authenticity

The Certificate of Authenticity (COA) label is used to:

- Upgrade the factory-installed operating system using the Microsoft Upgrade program for license validation.
- Reinstall the operating system because of a failure that has permanently disabled it.

The COA label location varies by server model. On rack-mounted server models, the COA label is located either on the front section of the right panel or on the right front corner of the top panel. On tower models, the COA label is located toward the rear of the top panel of the server.

8 Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

Related information

The following documents [and websites] provide related information:

- *HP StorageWorks P4000 G2 Unified NAS Gateway Quick Start Guide*
- HP StorageWorks P4000 SAN Solutions documentation
- http://www.hp.com/go/p4000

You can find these documents from the Manuals page of the HP Business Support Center website:

http://www.hp.com/support/manuals

In the Storage section, click **link label** and then select your product.

HP websites

For additional information, see the following HP websites:

- http://www.hp.com
- http://www.hp.com/go/storage
- http://www.hp.com/service_locator
- http://www.hp.com/support/manuals

- http://www.hp.com/support/downloads
- http://www.hp.com/storage/whitepapers

Typographic conventions

Table 6 Document conventions

Convention	Element
Blue text: Table 6	Cross-reference links and e-mail addresses
Blue, underlined text: http://www.hp.com	Website addresses
Bold text	• Keys that are pressed • Text typed into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
`Monospace` text	• File and directory names • System output • Code • Commands, their arguments, and argument values
`Monospace, italic` text	• Code variables • Command variables
`Monospace, bold` text	Emphasized monospace text

⚠ WARNING!

Indicates that failure to follow directions could result in bodily harm or death.

△ CAUTION:

Indicates that failure to follow directions could result in damage to equipment or data.

IMPORTANT:

Provides clarifying information or specific instructions.

NOTE:

Provides additional information.

TIP:

Provides helpful hints and shortcuts.

Rack stability

Rack stability protects personnel and equipment.

WARNING!

To reduce the risk of personal injury or damage to equipment:

- Extend leveling jacks to the floor.
- Ensure that the full weight of the rack rests on the leveling jacks.
- Install stabilizing feet on the rack.
- In multiple-rack installations, fasten racks together securely.
- Extend only one rack component at a time. Racks can become unstable if more than one component is extended.

Customer self repair

HP customer self repair (CSR) programs allow you to repair your StorageWorks product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider, or see the CSR website:

http://www.hp.com/go/selfrepair

9 System recovery

This chapter describes how to use the System Recovery DVD that is provided with your storage system.

The System Recovery DVD

The HP StorageWorks Storage System Recovery DVD that is provided with your storage system allows you to install an image or recover from a catastrophic failure.

At any later time, you may boot from the DVD and restore the server to the factory condition. This allows you to recover the system if all other means to boot the server fail.

While the recovery process makes every attempt to preserve the existing data volumes, you should have a backup of your data if at all possible before recovering the system.

As of HP StorageWorks Storage System Recovery DVD version 1.2, the DON'T ERASE volume is no longer used. If your system has a DON'T ERASE volume, the System Recovery process will ignore this volume.

NOTE:

Some storage systems do not include an internal DVD drive. For these systems, you must either use an external DVD drive to run the System Recovery DVD or create a USB Flash Drive that can then be used to complete the system recovery process. For more information, see Using a USB Flash Drive for System Recovery.

During the recovery process, the DVD overwrites the original OS logical drives. **All data on these drives is erased.**

Restore the factory image

1. Do one of the following:
   a. To use the direct connect access method, connect a keyboard, monitor, mouse, and DVD drive (if needed) directly to the server using a local I/O cable.
   b. To use the remote management access method, access the server using Integrated Lights-Out 2 (iLO 2) from a client PC.
2. Do one of the following:
   a. Insert the System Recovery DVD in the DVD drive.
   b. Insert the System Recovery DVD in the client PC.

3. Click **Restore Factory Image**.

   The upgrade process completes with little user intervention required. The server automatically reboots more than once.

IMPORTANT:

Do not interrupt the upgrade process.

When the upgrade process nears completion, the Windows Storage Server 2008 R2 desktop displays the following message: **The user's password must be changed before logging on the first time**. Log on to the storage system by establishing an Administrator password:

4. Click **OK**.
5. Type an Administrator password in the **New password** box.
6. Re-type the Administrator password in the **Confirm password** box.
7. Click the blue arrow next to the **Confirm password** box.
8. Click **OK**.

   After the Administrator password has been set, the storage system completes the upgrade process.
9. Remove the DVD or iLO 2 virtual DVD from the server.

Using a USB Flash Drive for System Recovery

Creating a System Recovery USB Flash drive is supported on Windows Vista, Windows 7, Windows Storage Server 2008, and Windows Storage Server 2008 R2 operating systems only.

If you create a backup copy of the System Recovery DVD using a USB Flash Drive, it can also be used to restore the system. To create system recovery media using a USB Flash drive follow the instructions below.

Create a System Recovery USB Flash Drive

1. Obtain a blank 8GB or larger USB Flash Drive.
2. Insert the USB Flash drive into your workstation or laptop.
3. Open an elevated command prompt with Administrator privileges.
4. At the command prompt, enter `diskpart`.
5. At the diskpart prompt, enter `list disk`.
6. Identify the disk number that corresponds to the flash drive. This is typically the last disk listed.
7. Enter `select disk <USB drive number>`. For example, `select disk 4`.
8. Enter `clean`.
9. Enter `create partition primary`.
10. Enter `select partition 1`.

11. Enter `format fs=fat32 quick`.

NOTE:

If your USB Flash Drive does not support the FAT32 file system, format the drive as NTFS instead. Omitting the `quick` parameter lengthens the format time considerably.

12. Enter `active` to mark the partition as active.
13. Enter `assign letter=<drive letter>` to assign a drive letter to the USB drive. For example, `assign letter=U`.
14. Insert the System Recovery DVD provided with the system.
15. Using Windows Explorer or a comparable utility, open the DVD so that all contents are visible.
16. Select all of the files (including bootmgr).
17. Copy all files to the root of the USB drive.

Use the USB Flash Drive for System Recovery

CAUTION:

During the recovery process, the System Recovery USB Flash drive overwrites the original OS logical drives. All data on these drives will be erased.

1. Do one of the following:
   a. To use the direct connect access method, connect a keyboard, monitor, and mouse, directly to the server using a local I/O cable.
   b. To use the remote management access method, access the server using Integrated Lights-Out 2 (iLO 2) from a client PC.
2. Do one of the following:
   a. Insert the System Recovery USB Flash drive in a USB port on the X Series system being restored.
   b. Insert the System Recovery USB Flash drive in the client PC connected to the iLO port of the X Series System being restored.
3. Click **Restore Factory Image**.

   The upgrade process completes with little user intervention required. The server automatically reboots more than once.

IMPORTANT:

Do not interrupt the upgrade process.

When the upgrade process nears completion, the Windows Storage Server 2008 R2 desktop displays the following message: **The user's password must be changed before logging on the first time**. Log on to the storage system by establishing an Administrator password:

4. Click **OK**.

5. Type an Administrator password in the **New password** box.
6. Re-type the Administrator password in the **Confirm password** box.
7. Click the blue arrow next to the **Confirm password** box.
8. Click **OK**.

   After the Administrator password has been set, the storage system completes the recovery process.
9. Remove the USB Flash drive from the X Series system or client PC.

Managing disks after a restoration

When a system that has existing data volumes (non operating system volumes) is restored using the System Recovery DVD, the data volumes will not have drive letters assigned to them. This is by design. The volume labels are retained and can be used to identify the data volumes. There is no workaround for this issue; however, drive letters can be assigned to volumes using **diskpart.exe** or by using Disk Management:

1. Click **Start > Run**, enter `diskmgmt.msc`and then click **OK**.
2. Right-click the disk and partition that you want to assign the drive letter to.
3. Select **Change drive Letter and Paths**.
4. In the **Change drive Letter and Paths** dialog box, select **Change**.
5. Select the appropriate drive letter, then click **OK**.
6. Click **Yes** to confirm the drive letter change.
7. Click **Yes** to continue. If the old drive letter needs to be re-used, reboot the server after clicking **Yes**.

A Regulatory compliance notices

This section contains regulatory notices for the HP StorageWorks family of products.

Regulatory compliance identification numbers

For the purpose of regulatory compliance certifications and identification, this product has been assigned a unique regulatory model number. The regulatory model number can be found on the product nameplate label, along with all required approval markings and information. When requesting compliance information for this product, always refer to this regulatory model number. The regulatory model number is not the marketing name or model number of the product.

Product specific information:

HP _____

Regulatory model number: _____

FCC and CISPR classification: _____

These products contain laser components. See Class 1 laser statement in the Laser compliance notices section.

Federal Communications Commission notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (for example, personal computers). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user.

FCC rating label

The FCC rating label on the device shows the classification (A or B) of the equipment. Class B devices have an FCC logo or ID on the label. Class A devices do not have an FCC logo or ID on the label. After you determine the class of the device, refer to the corresponding statement.

Class A equipment

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation

of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Class B equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit that is different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

Declaration of Conformity for products marked with the FCC logo, United States only

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions regarding this FCC declaration, contact us by mail or telephone:

- Hewlett-Packard Company P.O. Box 692000, Mail Stop 510101 Houston, Texas 77269-2000
- Or call 1-281-514-3333

Modification

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

When provided, connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC Rules and Regulations.

Canadian notice (Avis Canadien)

Class A equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la class A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Class B equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la class B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Union notice

Products bearing the CE marking comply with the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community.

Compliance with these directives implies conformity to the following European Norms (in parentheses are the equivalent international standards and regulations):

- EN 55022 (CISPR 22)—Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11)—Electromagnetic Immunity
- EN61000-3-2 (IEC61000-3-2)—Power Line Harmonics
- EN61000-3-3 (IEC61000-3-3)—Power Line Flicker
- EN 60950 (IEC60950)—Product Safety

Japanese notices

Japanese VCCI-A notice

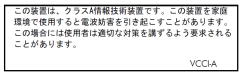

Japanese VCCI-B notice

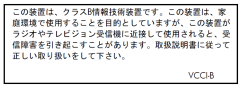

Japanese power cord statement

製品には、同梱された電源コードをお使い下さい。
同梱された電源コードは、他の製品では使用出来ません。

Please use the attached power cord.
The attached power cord is not allowed to use with other product.

Korean notices

Class A equipment

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Class B equipment

B급 기기 (가정용 정보통신기기)

이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거지역에서는 물론 모든지역에서 사용할 수 있습니다.

Taiwanese notices

BSMI Class A notice

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Taiwan battery recycle statement

Recovery mark: • Four-in-one recycling symbol	Recovery text: • "Please recycle waste batteries"
	廢電池請回收

Laser compliance notices

English laser notice

This device may contain a laser that is classified as a Class 1 Laser Product in accordance with U.S. FDA regulations and the IEC 60825-1. The product does not emit hazardous laser radiation.

⚠ WARNING!

Use of controls or adjustments or performance of procedures other than those specified herein or in the laser product's installation guide may result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- Do not try to open the module enclosure. There are no user-serviceable components inside.
- Do not operate controls, make adjustments, or perform procedures to the laser device other than those specified herein.
- Allow only HP Authorized Service technicians to repair the unit.

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1, 1976. Compliance is mandatory for products marketed in the United States.

Dutch laser notice

WAARSCHUWING: dit apparaat bevat mogelijk een laser die is geclassificeerd als een laserproduct van Klasse 1 overeenkomstig de bepalingen van de Amerikaanse FDA en de richtlijn IEC 60825-1. Dit product geeft geen gevaarlijke laserstraling af.

Als u bedieningselementen gebruikt, instellingen aanpast of procedures uitvoert op een andere manier dan in deze publicatie of in de installatiehandleiding van het laserproduct wordt aangegeven, loopt u het risico te worden blootgesteld aan gevaarlijke straling. Het risico van blootstelling aan gevaarlijke straling beperkt u als volgt:

- Probeer de behuizing van de module niet te openen. U mag zelf geen onderdelen repareren.
- Gebruik voor de laserapparatuur geen andere knoppen of instellingen en voer geen andere aanpassingen of procedures uit dan die in deze handleiding worden beschreven.
- Alleen door HP geautoriseerde technici mogen het apparaat repareren.

French laser notice

AVERTISSEMENT : cet appareil peut être équipé d'un laser classé en tant que Produit laser de classe 1 et conforme à la réglementation de la FDA américaine et à la norme 60825-1 de l'IEC. Ce produit n'émet pas de rayonnement dangereux.

L'utilisation de commandes, de réglages ou de procédures autres que ceux qui sont indiqués ici ou dans le manuel d'installation du produit laser peut exposer l'utilisateur à des rayonnements dangereux. Pour réduire le risque d'exposition à des rayonnements dangereux :

- Ne tentez pas d'ouvrir le boîtier renfermant l'appareil laser. Il ne contient aucune pièce dont la maintenance puisse être effectuée par l'utilisateur.
- Tout contrôle, réglage ou procédure autre que ceux décrits dans ce chapitre ne doivent pas être effectués par l'utilisateur.
- Seuls les Mainteneurs Agréés HP sont habilités à réparer l'appareil laser.

German laser notice

VORSICHT: Dieses Gerät enthält möglicherweise einen Laser, der nach den US-amerikanischen FDA-Bestimmungen und nach IEC 60825-1 als Laserprodukt der Klasse 1 zertifiziert ist. Gesundheitsschädliche Laserstrahlen werden nicht emittiert.

Die Anleitungen in diesem Dokument müssen befolgt werden. Bei Einstellungen oder Durchführung sonstiger Verfahren, die über die Anleitungen in diesem Dokument bzw. im Installationshandbuch des Lasergeräts hinausgehen, kann es zum Austritt gefährlicher Strahlung kommen. Zur Vermeidung der Freisetzung gefährlicher Strahlungen sind die folgenden Punkte zu beachten:

- Versuchen Sie nicht, die Abdeckung des Lasermoduls zu öffnen. Im Inneren befinden sich keine Komponenten, die vom Benutzer gewartet werden können.
- Benutzen Sie das Lasergerät ausschließlich gemäß den Anleitungen und Hinweisen in diesem Dokument.
- Lassen Sie das Gerät nur von einem HP Servicepartner reparieren.

Italian laser notice

AVVERTENZA: AVVERTENZA Questo dispositivo può contenere un laser classificato come prodotto laser di Classe 1 in conformità alle normative US FDA e IEC 60825-1. Questo prodotto non emette radiazioni laser pericolose.

L'eventuale esecuzione di comandi, regolazioni o procedure difformi a quanto specificato nella presente documentazione o nella guida di installazione del prodotto può causare l'esposizione a radiazioni nocive. Per ridurre i rischi di esposizione a radiazioni pericolose, attenersi alle seguenti precauzioni:

- Non cercare di aprire il contenitore del modulo. All'interno non vi sono componenti soggetti a manutenzione da parte dell'utente.
- Non eseguire operazioni di controllo, regolazione o di altro genere su un dispositivo laser ad eccezione di quelle specificate da queste istruzioni.
- Affidare gli interventi di riparazione dell'unità esclusivamente ai tecnici dell'Assistenza autorizzata HP.

Japanese laser notice

警告:本製品には、US FDA規則およびIEC 60825-1に基づくClass 1レーザー製品が含まれている場合があります。本製品は人体に危険なレーザー光は発しません。

本書およびレーザー製品のインストール ガイドに示されている以外の方法で制御、調整、使用した場合、人体に危険な光線にさらされる場合があります。人体に危険な光線にさらされないため、以下の項目を守ってください。

- モジュール エンクロージャを開けないでください。ユーザーが取り扱えるコンポーネントは含まれていません。
- 本書に示されている以外の方法で、レーザー デバイスを制御、調整、使用しないでください。
- HPの正規サービス技術者のみが本ユニットの修理を許可されています。

Spanish laser notice

ADVERTENCIA: Este dispositivo podría contener un láser clasificado como producto de láser de Clase 1 de acuerdo con la normativa de la FDA de EE.UU. e IEC 60825-1. El producto no emite radiaciones láser peligrosas.

El uso de controles, ajustes o manipulaciones distintos de los especificados aquí o en la guía de instalación del producto de láser puede producir una exposición peligrosa a las radiaciones. Para evitar el riesgo de exposición a radiaciones peligrosas:

- No intente abrir la cubierta del módulo. Dentro no hay componentes que el usuario pueda reparar.
- No realice más operaciones de control, ajustes o manipulaciones en el dispositivo láser que los aquí especificados.
- Sólo permita reparar la unidad a los agentes del servicio técnico autorizado HP.

Recycling notices

English notice

Disposal of Waste Equipment by Users in Private Households in the European Union

This symbol means do not dispose of your product with your other household waste. Instead, you should protect human health and the environment by handing over your waste equipment to a designated collection point for the recycling of waste electrical and electronic equipment. For more information, please contact your household waste disposal service

Bulgarian notice

Изхвърляне на отпадъчно оборудване от потребители в частни домакинства в Европейския съюз
Този символ върху продукта или опаковката му показва, че продуктът не трябва да се изхвърля заедно с другите битови отпадъци. Вместо това, трябва да предпазите човешкото здраве и околната среда, като предадете отпадъчното оборудване в предназначен за събирането му пункт за рециклиране на неизползваемо електрическо и електронно борудване. За допълнителна информация се свържете с фирмата по чистота, чиито услуги използвате.

Czech notice

Likvidace zařízení v domácnostech v Evropské unii
Tento symbol znamená, že nesmíte tento produkt likvidovat spolu s jiným domovním odpadem. Místo toho byste měli chránit lidské zdraví a životní prostředí tím, že jej předáte na k tomu určené sběrné pracoviště, kde se zabývají recyklací elektrického a elektronického vybavení. Pro více informací kontaktujte společnost zabývající se sběrem a svozem domovního odpadu.

Danish notice

Bortskaffelse af brugt udstyr hos brugere i private hjem i EU
Dette symbol betyder, at produktet ikke må bortskaffes sammen med andet husholdningsaffald. Du skal i stedet den menneskelige sundhed og miljøet ved at afl evere dit brugte udstyr på et dertil beregnet indsamlingssted for af brugt, elektrisk og elektronisk udstyr. Kontakt nærmeste renovationsafdeling for yderligere oplysninger.

Dutch notice

Inzameling van afgedankte apparatuur van particuliere huishoudens in de Europese Unie
Dit symbool betekent dat het product niet mag worden gedeponeerd bij het overige huishoudelijke afval. Bescherm de gezondheid en het milieu door afgedankte apparatuur in te leveren bij een hiervoor bestemd inzamelpunt voor recycling van afgedankte elektrische en elektronische apparatuur. Neem voor meer informatie contact op met uw gemeentereinigingsdienst.

Estonian notice

Äravisatavate seadmete likvideerimine Euroopa Liidu eramajapidamistes
See märk näitab, et seadet ei tohi visata olmeprügi hulka. Inimeste tervise ja keskkonna säästmise nimel tuleb äravisatav toode tuua elektriliste ja elektrooniliste seadmete käitlemisega egelevasse kogumispunkti. Küsimuste korral pöörduge kohaliku prügikäitlusettevõtte poole.

Finnish notice

Kotitalousjätteiden hävittäminen Euroopan unionin alueella
Tämä symboli merkitsee, että laitetta ei saa hävittää muiden kotitalousjätteiden mukana. Sen sijaan sinun on suojattava ihmisten terveyttä ja ympäristöä toimittamalla käytöstä poistettu laite sähkö- tai elektroniikkajätteen kierrätyspisteeseen. Lisätietoja saat jätehuoltoyhtiöltä.

French notice

Mise au rebut d'équipement par les utilisateurs privés dans l'Union Européenne
Ce symbole indique que vous ne devez pas jeter votre produit avec les ordures ménagères. Il est de votre responsabilité de protéger la santé et l'environnement et de vous débarrasser de votre équipement en le remettant à une déchetterie effectuant le recyclage des équipements électriques et électroniques. Pour de plus amples informations, prenez contact avec votre service d'élimination des ordures ménagères.

German notice

Entsorgung von Altgeräten von Benutzern in privaten Haushalten in der EU
Dieses Symbol besagt, dass dieses Produkt nicht mit dem Haushaltsmüll entsorgt werden darf. Zum Schutze der Gesundheit und der Umwelt sollten Sie stattdessen Ihre Altgeräte zur Entsorgung einer dafür vorgesehenen Recyclingstelle für elektrische und elektronische Geräte übergeben. Weitere Informationen erhalten Sie von Ihrem Entsorgungsunternehmen für Hausmüll.

Greek notice

Απόρριψη άχρηστου εξοπλισμού από ιδιώτες χρήστες στην Ευρωπαϊκή Ένωση

Αυτό το σύμβολο σημαίνει ότι δεν πρέπει να απορρίψετε το προϊόν με τα λοιπά οικιακά απορρίμματα. Αντίθετα, πρέπει να προστατέψετε την ανθρώπινη υγεία και το περιβάλλον παραδίδοντας τον άχρηστο εξοπλισμό σας σε εξουσιοδοτημένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Για περισσότερες πληροφορίες, επικοινωνήστε με την υπηρεσία απόρριψης απορριμμάτων της περιοχής σας.

Hungarian notice

A hulladék anyagok megsemmisítése az Európai Unió háztartásaiban

Ez a szimbólum azt jelzi, hogy a készüléket nem szabad a háztartási hulladékkal együtt kidobni. Ehelyett a leselejtezett berendezéseknek az elektromos vagy elektronikus hulladék átvételére kijelölt helyen történő beszolgáltatásával megóvja az emberi egészséget és a környezetet.További információt a helyi köztisztasági vállalattól kaphat.

Italian notice

Smaltimento di apparecchiature usate da parte di utenti privati nell'Unione Europea

Questo simbolo avvisa di non smaltire il prodotto con i normali rifi uti domestici. Rispettare la salute umana e l'ambiente conferendo l'apparecchiatura dismessa a un centro di raccolta designato per il riciclo di apparecchiature elettroniche ed elettriche. Per ulteriori informazioni, rivolgersi al servizio per lo smaltimento dei rifi uti domestici.

Latvian notice

Europos Sąjungos namų ūkio vartotojų įrangos atliekų šalinimas

Šis simbolis nurodo, kad gaminio negalima išmesti kartu su kitomis buitinėmis atliekomis. Kad apsaugotumėte žmonių sveikatą ir aplinką, pasenusią nenaudojamą įrangą turite nuvežti į elektrinių ir elektroninių atliekų surinkimo punktą. Daugiau informacijos teiraukitės buitinių atliekų surinkimo tarnybos.

Lithuanian notice

Nolietotu iekārtu iznīcināšanas noteikumi lietotājiem Eiropas Savienības privātajās mājsaimniecībās
Šis simbols norāda, ka ierīci nedrīkst utilizēt kopā ar citiem mājsaimniecības atkritumiem. Jums jārūpējas par cilvēku veselības un vides aizsardzību, nododot lietoto aprīkojumu otrreizējai pārstrādei īpašā lietotu elektrisko un elektronisko ierīču savākšanas punktā. Lai iegūtu plašāku informāciju, lūdzu, sazinieties ar savu mājsaimniecības atkritumu likvidēšanas dienestu.

Polish notice

Utylizacja zużytego sprzętu przez użytkowników w prywatnych gospodarstwach domowych
w krajach Unii Europejskiej
Ten symbol oznacza, że nie wolno wyrzucać produktu wraz z innymi domowymi odpadkami. Obowiązkiem użytkownika jest ochrona zdrowa ludzkiego i środowiska przez przekazanie zużytego sprzętu do wyznaczonego punktu zajmującego się recyklingiem odpadów powstałych ze sprzętu elektrycznego i elektronicznego. Więcej informacji można uzyskać od lokalnej firmy zajmującej wywozem nieczystości.

Portuguese notice

Descarte de equipamentos usados por utilizadores domésticos na União Europeia
Este símbolo indica que não deve descartar o seu produto juntamente com os outros lixos domiciliares. Ao invés disso, deve proteger a saúde humana e o meio ambiente levando o seu equipamento para descarte em um ponto de recolha destinado à reciclagem de resíduos de equipamentos eléctricos e electrónicos. Para obter mais informações, contacte o seu serviço de tratamento de resíduos domésticos.

Romanian notice

Casarea echipamentului uzat de către utilizatorii casnici din Uniunea Europeană
Acest simbol înseamnă să nu se arunce produsul cu alte deşeuri menajere. În schimb, trebuie să protejaţi sănătatea umană şi mediul predând echipamentul uzat la un punct de colectare desemnat pentru reciclarea echipamentelor electrice şi electronice uzate. Pentru informaţii suplimentare, vă rugăm să contactaţi serviciul de eliminare a deşeurilor menajere local.

Slovak notice

Likvidácia vyradených zariadení používateľmi v domácnostiach v Európskej únii

Tento symbol znamená, že tento produkt sa nemá likvidovať s ostatným domovým odpadom. Namiesto toho by ste mali chrániť ľudské zdravie a životné prostredie odovzdaním odpadového zariadenia na zbernom mieste, ktoré je určené na recykláciu odpadových elektrických a elektronických zariadení. Ďalšie informácie získate od spoločnosti zaoberajúcej sa likvidáciou domového odpadu.

Spanish notice

Eliminación de los equipos que ya no se utilizan en entornos domésticos de la Unión Europea

Este símbolo indica que este producto no debe eliminarse con los residuos domésticos. En lugar de ello, debe evitar causar daños a la salud de las personas y al medio ambiente llevando los equipos que no utilice a un punto de recogida designado para el reciclaje de equipos eléctricos y electrónicos que ya no se utilizan. Para obtener más información, póngase en contacto con el servicio de recogida de residuos domésticos.

Swedish notice

Hantering av elektroniskt avfall för hemanvändare inom EU

Den här symbolen innebär att du inte ska kasta din produkt i hushållsavfallet. Värna i stället om natur och miljö genom att lämna in uttjänt utrustning på anvisad insamlingsplats. Allt elektriskt och elektroniskt avfall går sedan vidare till återvinning. Kontakta ditt återvinningsföretag för mer information.

Turkish notice

Türkiye Cumhuriyeti: EEE Yönetmeliğine Uygundur

Battery replacement notices

Dutch battery notice

Verklaring betreffende de batterij

WAARSCHUWING: dit apparaat bevat mogelijk een batterij.

- Probeer de batterijen na het verwijderen niet op te laden.
- Stel de batterijen niet bloot aan water of temperaturen boven 60° C.
- De batterijen mogen niet worden beschadigd, gedemonteerd, geplet of doorboord.
- Zorg dat u geen kortsluiting veroorzaakt tussen de externe contactpunten en laat de batterijen niet in aanraking komen met water of vuur.
- Gebruik ter vervanging alleen door HP goedgekeurde batterijen.

Batterijen, accu's en accumulators mogen niet worden gedeponeerd bij het normale huishoudelijke afval. Als u de batterijen/accu's wilt inleveren voor hergebruik of op de juiste manier wilt vernietigen, kunt u gebruik maken van het openbare inzamelingssysteem voor klein chemisch afval of ze terugsturen naar HP of een geautoriseerde HP Business of Service Partner.

Neem contact op met een geautoriseerde leverancier of een Business of Service Partner voor meer informatie over het vervangen of op de juiste manier vernietigen van accu's.

French battery notice

Avis relatif aux piles

AVERTISSEMENT : cet appareil peut contenir des piles.

- N'essayez pas de recharger les piles après les avoir retirées.
- Évitez de les mettre en contact avec de l'eau ou de les soumettre à des températures supérieures à 60°C.
- N'essayez pas de démonter, d'écraser ou de percer les piles.
- N'essayez pas de court-circuiter les bornes de la pile ou de jeter cette dernière dans le feu ou l'eau.
- Remplacez les piles exclusivement par des pièces de rechange HP prévues pour ce produit.

Les piles, modules de batteries et accumulateurs ne doivent pas être jetés avec les déchets ménagers. Pour permettre leur recyclage ou leur élimination, veuillez utiliser les systèmes de collecte publique ou renvoyez-les à HP, à votre Partenaire Agréé HP ou aux agents agréés.

Contactez un Revendeur Agréé ou Mainteneur Agréé pour savoir comment remplacer et jeter vos piles.

German battery notice

Hinweise zu Batterien und Akkus

VORSICHT: Dieses Produkt enthält unter Umständen eine Batterie oder einen Akku.

- Versuchen Sie nicht, Batterien und Akkus außerhalb des Gerätes wieder aufzuladen.
- Schützen Sie Batterien und Akkus vor Feuchtigkeit und Temperaturen über 60°.
- Verwenden Sie Batterien und Akkus nicht missbräuchlich, nehmen Sie sie nicht auseinander und vermeiden Sie mechanische Beschädigungen jeglicher Art.
- Vermeiden Sie Kurzschlüsse, und setzen Sie Batterien und Akkus weder Wasser noch Feuer aus.
- Ersetzen Sie Batterien und Akkus nur durch die von HP vorgesehenen Ersatzteile.

Batterien und Akkus dürfen nicht über den normalen Hausmüll entsorgt werden. Um sie der Wiederverwertung oder dem Sondermüll zuzuführen, nutzen Sie die öffentlichen Sammelstellen, oder setzen Sie sich bezüglich der Entsorgung mit einem HP Partner in Verbindung.

Weitere Informationen zum Austausch von Batterien und Akkus oder zur sachgemäßen Entsorgung erhalten Sie bei Ihrem HP Partner oder Servicepartner.

Italian battery notice

Istruzioni per la batteria

AVVERTENZA: Questo dispositivo può contenere una batteria.

- Non tentare di ricaricare le batterie se rimosse.
- Evitare che le batterie entrino in contatto con l'acqua o siano esposte a temperature superiori a 60° C.
- Non smontare, schiacciare, forare o utilizzare in modo improprio la batteria.
- Non accorciare i contatti esterni o gettare in acqua o sul fuoco la batteria.
- Sostituire la batteria solo con i ricambi HP previsti a questo scopo.

Le batterie e gli accumulatori non devono essere smaltiti insieme ai rifiuti domestici. Per procedere al riciclaggio o al corretto smaltimento, utilizzare il sistema di raccolta pubblico dei rifiuti o restituirli a HP, ai Partner Ufficiali HP o ai relativi rappresentanti.

Per ulteriori informazioni sulla sostituzione e sullo smaltimento delle batterie, contattare un Partner Ufficiale o un Centro di assistenza autorizzato.

Japanese battery notice

バッテリに関する注意

警告:本製品はバッテリを内蔵している場合があります。

- バッテリを取り外している場合は、充電しないでください。
- バッテリを水にさらしたり、60°C (140°F) 以上の温度にさらさないでください。
- バッテリを誤用、分解、破壊したり、穴をあけたりしないでください。
- 外部極を短絡させたり、火や水に投棄しないでください。
- バッテリを交換する際は、HP指定の製品と交換してください。

バッテリ、バッテリ パック、蓄電池は一般の家庭廃棄物と一緒に廃棄しないでください。リサイクルまたは適切に廃棄するため、公共の収集システム、HP、HPパートナー、またはHPパートナーの代理店にお送りください。

バッテリ交換および適切な廃棄方法についての情報は、HPのサポート窓口にお問い合わせください。

Spanish battery notice

Declaración sobre las baterías

ADVERTENCIA: Este dispositivo podría contener una batería.

- No intente recargar las baterías si las extrae.
- Evite el contacto de las baterías con agua y no las exponga a temperaturas superiores a los 60 °C (140 °F).
- No utilice incorrectamente, ni desmonte, aplaste o pinche las baterías.
- No cortocircuite los contactos externos ni la arroje al fuego o al agua.
- Sustituya las baterías sólo por el repuesto designado por HP.

Las baterías, los paquetes de baterías y los acumuladores no se deben eliminar junto con los desperdicios generales de la casa. Con el fin de tirarlos al contenedor de reciclaje adecuado, utilice los sistemas públicos de recogida o devuélvalas a HP, un distribuidor autorizado de HP o sus agentes.

Para obtener más información sobre la sustitución de la batería o su eliminación correcta, consulte con su distribuidor o servicio técnico autorizado.

Index

G

H

J

K

L

M

N

P

Q

R

S

T

U

V

W